Nested Regular Path Queries in Description Logics

Meghyn Bienvenu
Lab. de Recherche en Informatique
CNRS & Univ. Paris Sud, France

Diego Calvanese
KRDB Research Centre
Free Univ. of Bozen-Bolzano, Italy

Magdalena Ortiz
Mantas Šimkus
Institute of Information Systems
Vienna Univ. of Technology, Austria

arXiv:1402.7122v2 [cs.LO] 4 Mar 2014

Abstract

Two-way regular path queries (2RPQs) have received increased attention recently due to their ability to relate pairs of objects by flexibly navigating graph-structured data. They are present in property paths in SPARQL 1.1, the new standard RDF query language, and in the XML query language XPath. In line with XPath, we consider the extension of 2RPQs with nesting, which allows one to require that objects along a path satisfy complex conditions, in turn expressed through (nested) 2RPQs. We study the computational complexity of answering nested 2RPQs and conjunctions thereof (CN2RPQs) in the presence of domain knowledge expressed in description logics (DLs). We establish tight complexity bounds in data and combined complexity for a variety of DLs, ranging from lightweight DLs (*DL-Lite*, $\mathcal{EL}$) up to highly expressive ones. Interestingly, we are able to show that adding nesting to (C)2RPQs does not affect worst-case data complexity of query answering for any of the considered DLs. However, in the case of lightweight DLs, adding nesting to 2RPQs leads to a surprising jump in combined complexity, from P-complete to EXP-complete.

1 Introduction

Both in knowledge representation and in databases, there has been great interest recently in expressive mechanisms for querying data, while taking into account complex domain knowledge (Calvanese, De Giacomo, and Lenzerini 2008; Glimm et al. 2008). Description Logics (DLs) (Baader et al. 2003), which on the one hand underlie the W3C standard Web Ontology Language (OWL), and on the other hand are able to capture at the intensional level conceptual modeling formalisms like UML and ER, are considered particularly well suited for representing a domain of interest (Borgida and Brachman 2003). In DLs, instance data, stored in a so-called ABox, is constituted by ground facts over unary and binary predicates (*concepts* and *roles*, respectively), and hence resembles data stored in graph databases (Consens and Mendelzon 1990; Barceló et al. 2012). There is a crucial difference, however, between answering queries over graph databases and over DL ABoxes. In the former, the data is assumed to be complete, hence query answering amounts to the standard database task of query evaluation. In the latter, it is typically assumed that the data is incomplete and additional domain knowledge is provided by the DL ontology (or TBox). Hence query answering amounts to the more complex task of computing *certain answers*, i.e., those answers that are obtained from all databases that both contain the explicit facts in the ABox and satisfy the TBox constraints. This difference has driven research in different directions.

In databases, expressive query languages for querying graph-structured data have been studied, which are based on the requirement of relating objects by flexibly navigating the data. The main querying mechanism that has been considered for this purpose is that of one-way and two-way regular path queries (RPQs and 2RPQs) (Cruz, Mendelzon, and Wood 1987; Calvanese et al. 2003), which are queries returning pairs of objects related by a *path* whose sequence of edge labels belongs to a regular language over the (binary) database relations and their inverses. Conjunctive 2RPQs (C2RPQs) (Calvanese et al. 2000) are a significant extension of such queries that add to the navigational ability the possibility of expressing arbitrary selections, projections, and joins over objects related by 2RPQs, in line with conjunctive queries (CQs) over relational databases. Two-way RPQs are present in the property paths in SPARQL 1.1 (Harris and Seaborne 2013), the new standard RDF query language, and in the XML query language XPath (Berglund and others 2010). An additional construct that is present in XPath is the possibility of using *existential test operators*, also known as *nesting*, to express sophisticated conditions along navigation paths. When an existential test $\langle E \rangle$ is used in a 2RPQ E', there will be objects along the main navigation path for E' that match positions of E' where $\langle E \rangle$ appears; such objects are required to be the origin of a path conforming to the (nested) 2RPQ E. It is important to notice that existential tests in general cannot be captured even by C2RPQs, e.g., when tests appear within a transitive closure of an RPQ. Hence, adding nesting effectively increases the expressive power of 2RPQs and of C2RPQs.

In the DL community, query answering has been investigated extensively for a wide range of DLs, with much of the work devoted to CQs. With regards to the complexity of query answering, attention has been paid on the one hand to *combined complexity*, i.e., the complexity measured considering as input both the query and the DL knowledge base (constituted by TBox and ABox), and on the other

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

hand to *data complexity*, i.e., when only the ABox is considered as input. For expressive DLs that extend $\mathcal{ALC}$, CQ answering is typically coNP-complete in data-complexity (Ortiz, Calvanese, and Eiter 2008), and 2EXP-complete in combined complexity (Glimm et al. 2008; Lutz 2008; Eiter et al. 2009). For lightweight DLs, instead, CQ answering is in AC^0 in data complexity for *DL-Lite* (Calvanese et al. 2007), and P-complete for $\mathcal{EL}$ (Krisnadhi and Lutz 2007). For both logics, the combined complexity is dominated by the NP-completeness of CQ evaluation over plain relational databases. There has also been some work on (2)RPQs and C(2)RPQs. For the very expressive DLs $\mathcal{ZIQ}$, $\mathcal{ZOQ}$, and $\mathcal{ZOI}$, where regular expressions over roles are present also in the DL, a 2EXP upper bound has been shown via techniques based on alternating automata over infinite trees (Calvanese, Eiter, and Ortiz 2009). For the Horn fragments of $\mathcal{SHOIQ}$ and $\mathcal{SROIQ}$, P-completeness in data complexity and EXP/2EXP-completeness in combined complexity are known (Ortiz, Rudolph, and Simkus 2011). For lightweight DLs, tight bounds for answering 2RPQs and C2RPQs have only very recently been established by Bienvenu, Ortiz, and Simkus (2013): for (C)(2)RPQs, data complexity is NL-complete in *DL-Lite* and *DL-Lite*$_\mathcal{R}$, and P-complete in $\mathcal{EL}$ and $\mathcal{ELH}$. For all of these logics, combined complexity is P-complete for (2)RPQs and PSPACE-complete for C(2)RPQs.

Motivated by the expressive power of nesting in XPath and SPARQL, in this paper we significantly advance these latter lines of research on query answering in DLs, and study the impact of adding nesting to 2RPQs and C2RPQs. We establish tight complexity bounds in data and combined complexity for a variety of DLs, ranging from lightweight DLs of the *DL-Lite* and $\mathcal{EL}$ families up to the highly expressive ones of the $\mathcal{SH}$ and $\mathcal{Z}$ families. Our results are summarized in Table 1. For DLs containing at least $\mathcal{ELI}$, we are able to encode away nesting, thus showing that the worst-case complexity of query answering is not affected by this construct. Instead, for lightweight DLs (starting already from *DL-Lite*!), we show that adding nesting to 2RPQs leads to a surprising jump in combined complexity, from P-complete to EXP-complete. We then develop a sophisticated rewriting-based technique that builds on (but significantly extends) the one proposed by Bienvenu, Ortiz, and Simkus (2013), which we use to prove that the problem remains in NL for *DL-Lite*. We thus show that adding nesting to (C)2RPQs does not affect worst-case data complexity of query answering for lightweight DLs.

Some proofs have been relegated to the appendix.

2 Preliminaries

We briefly recall the syntax and semantics of description logics (DLs). As usual, we assume countably infinite, mutually disjoint sets $\mathsf{N_C}$, $\mathsf{N_R}$, and $\mathsf{N_I}$ of *concept names*, *role names*, and *individuals*. We typically use A for concept names, p for role names, and a, b for individuals. An *inverse role* takes the form p^- where $p \in \mathsf{N_R}$. We let $\mathsf{N_R^\pm} = \mathsf{N_R} \cup \{p^- \mid p \in \mathsf{N_R}\}$ and denote by r elements of $\mathsf{N_R^\pm}$.

A DL knowledge base (KB) consists of a *TBox* and an *ABox*, whose forms depend on the DL in question. In the DL $\mathcal{ELHI}_\bot$, a TBox is defined as a set of *(positive) role inclusions* of the form $r \sqsubseteq r'$ and *negative role inclusions* of the form $r \sqcap r' \sqsubseteq \bot$ with $r, r' \in \mathsf{N_R^\pm}$, and *concept inclusions* of the form $C \sqsubseteq D$, where C and D are *complex concepts* formed according to the following syntax:[1]

$$C ::= \top \mid \bot \mid A \mid \exists r.C \mid C \sqcap C$$

with $A \in \mathsf{N_C}$ and $r \in \mathsf{N_R^\pm}$.

Some of our results refer specifically to the *lightweight DLs* that we define next. $\mathcal{ELHI}$ is the fragment of $\mathcal{ELHI}_\bot$ that has no $\bot$. $\mathcal{ELH}$ and $\mathcal{ELI}$ are obtained by additionally disallowing inverse roles and role inclusions, respectively. *DL-Lite*$_\mathcal{R}$ is also a fragment of $\mathcal{ELHI}_\bot$, in which concept inclusions can only take the forms $B_1 \sqsubseteq B_2$ and $B_1 \sqcap B_2 \sqsubseteq \bot$, for B_i a concept name or concept of the form $\exists r.\top$ with $r \in \mathsf{N_R^\pm}$. *DL-Lite* is the fragment of *DL-Lite*$_\mathcal{R}$ that disallows (positive and negative) role inclusions.

An *ABox* is a set of assertions of the form $C(a)$ or $r(a, b)$, where C is a complex concept, $r \in \mathsf{N_R^\pm}$, and $a, b \in \mathsf{N_I}$. We use $\mathsf{Ind}(\mathcal{A})$ to refer to the set of individuals in $\mathcal{A}$.

Semantics. The semantics of DL KBs is based upon *interpretations*, which take the form $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty set and $\cdot^\mathcal{I}$ maps each $a \in \mathsf{N_I}$ to $a^\mathcal{I} \in \Delta^\mathcal{I}$, each $A \in \mathsf{N_C}$ to $A^\mathcal{I} \subseteq \Delta^\mathcal{I}$, and each $p \in \mathsf{N_R}$ to $p^\mathcal{I} \subseteq \Delta^\mathcal{I} \times \Delta^\mathcal{I}$.[2] The function $\cdot^\mathcal{I}$ can be straightforwardly extended to complex concepts and roles. In the case of $\mathcal{ELHI}_\bot$, this is done as follows: $\top^\mathcal{I} = \Delta^\mathcal{I}$, $\bot^\mathcal{I} = \emptyset$, $(p^-)^\mathcal{I} = \{(c, d) \mid (d, c) \in p^\mathcal{I}\}$, $(\exists r.C)^\mathcal{I} = \{c \mid \exists d : (c, d) \in r^\mathcal{I}, d \in C^\mathcal{I}\}$, and $(C \sqcap D)^\mathcal{I} = C^\mathcal{I} \cap D^\mathcal{I}$. An interpretation $\mathcal{I}$ satisfies an inclusion $G \sqsubseteq H$ if $G^\mathcal{I} \subseteq H^\mathcal{I}$, and it satisfies an assertion $C(a)$ (resp., $r(a, b)$) if $a^\mathcal{I} \in A^\mathcal{I}$ (resp., $(a^\mathcal{I}, b^\mathcal{I}) \in r^\mathcal{I}$). A *model* of a KB $(\mathcal{T}, \mathcal{A})$ is an interpretation $\mathcal{I}$ which satisfies all inclusions in $\mathcal{T}$ and assertions in $\mathcal{A}$.

Complexity. In addition to P and (co)NP, our results refer to the complexity classes NL (non-deterministic logarithmic space), PSPACE (polynomial space), and (2)EXP ((double) exponential time), cf. (Papadimitriou 1993).

3 Nested Regular Path Queries

We now introduce our query languages. In RPQs, nested RPQs and their extensions, *atoms* are given by (nested) regular expressions whose symbols are *roles*. The set Roles of roles contains $\mathsf{N_R^\pm}$, and all *test roles* of the forms $\{a\}?$ and $A?$ with $a \in \mathsf{N_I}$ and $A \in \mathsf{N_C}$. They are interpreted as $(\{a\}?)^\mathcal{I} = (a^\mathcal{I}, a^\mathcal{I})$, and $(A?)^\mathcal{I} = \{(o, o) \mid o \in A^\mathcal{I}\}$.

Definition 3.1. *A* nested regular expression *(NRE), denoted by E, is constructed according to the following syntax:*

$$E ::= \sigma \mid E \cdot E \mid E \cup E \mid E^* \mid \langle E \rangle$$

where $\sigma \in \mathsf{Roles}$.

[1] We slightly generalize the usual $\mathcal{ELHI}_\bot$ by allowing for negative role inclusions.

[2] Note that we do not make the *unique name assumption* (UNA), but all of our results continue to hold if the UNA is adopted.

	2RPQ		C2RPQ		N2RPQ / CN2RPQ	
	data	combined	data	combined	data	combined
Graph DBs & RDFS	NL-c	NL-c	NL-c	NP-c	NL-c	**P-c** / NP-c
DL-Lite	NL-c	P-c	NL-c	PSPACE-c	**NL-c**	**EXP-c**
Horn DLs (e.g., $\mathcal{EL}$, Horn-$\mathcal{SHIQ}$)	P-c	P-c	P-c	PSPACE-c	**P-c**	**EXP-c**
Expressive DLs (e.g., $\mathcal{ALC}$, $\mathcal{SHIQ}$)	coNP-h	EXP-c	coNP-h	2EXP-c	coNP-h	**EXP-c / 2EXP-c**

Table 1: Complexity of query answering. The 'c' indicates completeness, the 'h' hardness. New results are marked in bold. For existing results, refer to (Bienvenu, Ortiz, and Simkus 2013; Pérez, Arenas, and Gutierrez 2010; Barceló Baeza 2013; Calvanese, Eiter, and Ortiz 2009; Ortiz, Rudolph, and Simkus 2011) and references therein.

We assume a countably infinite set $\mathsf{N_V}$ *of* variables *(disjoint from* $\mathsf{N_C}$*,* $\mathsf{N_R}$*, and* $\mathsf{N_I}$*). Each* $t \in \mathsf{N_V} \cup \mathsf{N_I}$ *is a* term*. An* atom *is either a* concept atom *of the form* $A(t)$*, with* $A \in \mathsf{N_C}$ *and* t *a term, or a* role atom *of the form* $E(t, t')$*, with* E *an NRE and* t, t' *two (possibly equal) terms.*

A nested two-way regular path query *(N2RPQ)* $q(x, y)$ *is an atom of the form* $E(x, y)$*, where* E *is an NRE and* x, y *are two distinct variables. A* conjunctive N2RPQ *(CN2RPQ)* $q(\vec{x})$ *with* answer variables $\vec{x}$ *has the form* $\exists \vec{y}.\varphi$*, where* φ *is a conjunction of atoms whose variables are among* $\vec{x} \cup \vec{y}$*.*

A (plain) regular expression (RE) is an NRE that does not have subexpressions of the form $\langle E \rangle$. Two-way regular path queries *(2RPQs) and* conjunctive 2RPQs *(C2RPQs) are defined analogously to N2RPQs and CN2RPQs but allowing only plain REs in atoms.*

Given an interpretation $\mathcal{I}$, the semantics of an NRE E is defined by induction on its structure:

$$(E_1 \cdot E_2)^{\mathcal{I}} = E_1^{\mathcal{I}} \circ E_2^{\mathcal{I}},$$
$$(E_1 \cup E_2)^{\mathcal{I}} = E_1^{\mathcal{I}} \cup E_2^{\mathcal{I}},$$
$$(E_1^*)^{\mathcal{I}} = (E_1^{\mathcal{I}})^*,$$
$$\langle E \rangle^{\mathcal{I}} = \{(o, o) \mid \text{there is } o' \in \Delta^{\mathcal{I}} \text{ s.t. } (o, o') \in E^{\mathcal{I}}\}.$$

A *match* for a C2NRPQ $q(\vec{x}) = \exists \vec{y}.\varphi$ in an interpretation $\mathcal{I}$ is a mapping from the terms in φ to $\Delta^{\mathcal{I}}$ such that (i) $\pi(a) = a^{\mathcal{I}}$ for every individual a of φ, (ii) $\pi(x) \in A^{\mathcal{I}}$ for every concept atom $A(x)$ of φ, and (iii) $(\pi(x), \pi(y)) \in E^{\mathcal{I}}$ for every role atom $E(x, y)$ of φ. Let $\mathsf{ans}(q, \mathcal{I}) = \{\pi(\vec{x}) \mid \pi \text{ is a match for } q \text{ in } \mathcal{I}\}$. An individual tuple $\vec{a}$ with the same arity as $\vec{x}$ is called a *certain answer* to q over a KB $\langle \mathcal{T}, \mathcal{A} \rangle$ if $(\vec{a})^{\mathcal{I}} \in \mathsf{ans}(q, \mathcal{I})$ for every model $\mathcal{I}$ of $\langle \mathcal{T}, \mathcal{A} \rangle$. We use $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ to denote the set of all certain answers to q over $\langle \mathcal{T}, \mathcal{A} \rangle$. In what follows, by *query answering*, we will mean the problem of deciding whether $\vec{a} \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$.

Example 3.1. We consider an ABox of advisor relationships of PhD holders[3]. We assume an *advisor* relation between nodes representing academics. There are also nodes for theses, universities, research topics, and countries, related in the natural way via roles *wrote*, *subm*(itted), *topic*, and *loc*(ation). We give two queries over this ABox.

$$q_1(x, y) = (advisor \cdot \langle wrote \cdot topic \cdot Physics? \rangle)^* (x, y)$$

[3]Our examples are inspired by the *Mathematics Genealogy Project* (`http://genealogy.math.ndsu.nodak.edu/`).

Query q_1 is an N2RPQ that retrieves pairs of a person x and an academic ancestor y of x such that all people on the path from x to y (including y itself) wrote a thesis in Physics.

$$q_2(x, y, z) = advisor^-(x, z), \quad advisor^*(x, w),$$
$$advisor^- \cdot \langle wrote \cdot \langle topic \cdot DBs? \rangle \cdot subm \cdot loc \cdot \{usa\}? \rangle (y, z),$$
$$\left(advisor \cdot \langle wrote \cdot \langle topic \cdot Logic? \rangle \cdot subm \cdot loc \cdot EU? \rangle\right)^* (y, w)$$

Query q_2 is a CN2RPQ that looks for triples of individuals x, y, z such that x and y have both supervised z, who wrote a thesis on Databases and who submitted this thesis to a university in the USA. Moreover, x and y have a common ancestor w, and all people on the path from x to w, including w, must have written a thesis in Logic and must have submitted this thesis to a university in an EU country. ■

It will often be more convenient to deal with an automata-based representation of (C)N2RPQs, which we provide next.

Definition 3.2. *A* nested NFA *(n-NFA) has the form* $(\mathbf{A}, s_0, F_0)$ *where* $\mathbf{A}$ *is an indexed set* $\{\alpha_1, \ldots, \alpha_n\}$*, where each* $\alpha_l \in \mathbf{A}$ *is an automaton of the form* (S, s, δ, F)*, where* S *is a set of states,* $s \in S$ *is the initial state,* $F \subseteq S$ *is the set of final states, and*

$$\delta \subseteq S \times (\mathsf{Roles} \cup \{\langle j_1, \ldots, j_k \rangle \mid l < j_i \leq n, \text{ for } i \in \{1, \ldots, k\}\}) \times S$$

We assume that the sets of states of the automata in $\mathbf{A}$ *are pairwise disjoint, and we require that* $\{s_0\} \cup F_0$ *are states of a single automaton in* $\mathbf{A}$*. If in each transition* $(s, \langle j_1, \ldots, j_k \rangle, s')$ *of each automaton in* $\mathbf{A}$ *we have* $k = 1$*, then the n-NFA is called* reduced*.*

When convenient notationally, we will denote an *n*-NFA $(\mathbf{A}, s_0, F_0)$ by $\mathbf{A}_{s_0, F_0}$. Moreover, we will use S_i, δ_i, and F_i to refer to the states, transition relation, and final states of α_i.

Definition 3.3. *Given an interpretation* $\mathcal{I}$*, we define* $\mathbf{A}^{\mathcal{I}}_{s_0, F_0}$ *inductively as follows. Let* α_l *be the (unique) automaton in* $\mathbf{A}$ *such that* $\{s_0\} \cup F_0 \subseteq S_l$*. Then* $(o, o') \in \mathbf{A}^{\mathcal{I}}_{s_0, F_0}$ *if there is a sequence* $s_0 o_0 s_1 \cdots o_{k-1} s_k o_k$*, for* $k \geq 0$*, such that* $o_0 = o$*,* $o_k = o'$*,* $s_k \in F_0$*, and for* $i \in \{1, \ldots, k\}$ *there is a transition* $(s_{i-1}, \sigma_i, s_i) \in \delta_l$ *such that either*

- $\sigma_i \in \mathsf{Roles}$ *and* $(o_{i-1}, o_i) \in \sigma_i^{\mathcal{I}}$*, or*
- $\sigma_i = \langle j_1, \ldots, j_k \rangle$ *such that, for every* $m \in \{1, \ldots, k\}$*, there exists* $o_m \in \Delta^{\mathcal{I}}$ *with* $(o_i, o_m) \in \mathbf{A}^{\mathcal{I}}_{s', F'}$*, where* s' *and* F' *are the initial and final states of* α_{j_m} *respectively.*

Note that an n-NFA $\mathbf{A}_{s_0,F_0}$ such that there are no transitions of the form $(s, \langle j_1, \ldots, j_k \rangle, s')$ in the unique α_l with $\{s_0\} \cup F_0 \subseteq S_l$ is equivalent to a standard NFA.

For every NRE E one can construct in polynomial time an n-NFA $\mathbf{A}_{s_0,F_0}$ such that $E^{\mathcal{I}} = \mathbf{A}^{\mathcal{I}}_{s_0,F_0}$ for every interpretation $\mathcal{I}$. This is an almost immediate consequence of the correspondence between regular expressions and finite state automata. Moreover, any n-NFA can be transformed into an equivalent reduced n-NFA by introducing linearly many additional states. In the following, unless stated otherwise, we assume all n-NFAs are reduced.

4 Upper Bounds via Reductions

In this section, we derive some upper bounds on the complexity of answering (C)N2RPQs in different DLs, by means of reductions to other problems. For simplicity, we assume in the rest of this section that query atoms do not employ test roles of the form $\{a\}?$. This is without loss of generality, since each symbol $\{a\}?$ can be replaced by $A_a?$ for a fresh concept name A_a, by adding the ABox assertion $A_a(a)$.

We start by showing that answering CN2RPQs can be polynomially reduced to answering non-nested C2RPQs using TBox axioms that employ inverses, conjunction on the left, and qualified existential restrictions.

Proposition 4.1. *For each CN2RPQ q, one can compute in polynomial time an $\mathcal{ELI}$ TBox $\mathcal{T}'$ and C2RPQ q' such that $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle)$ for every KB $\langle \mathcal{T}, \mathcal{A} \rangle$.*

Proof. Let q be an arbitrary CN2RPQ whose role atoms are given by n-NFAs, that is, they take the form $\mathbf{A}_{s_0,F_0}(x, y)$.

For each atom $\mathbf{A}_{s_0,F_0}(x, y)$ in q and each $\alpha_i \in \mathbf{A}$, we use a fresh concept name A_s for each state $s \in S_i$, and define a TBox $\mathcal{T}_{\alpha_i}$ that contains:

- $\top \sqsubseteq A_f$ for each $f \in F_i$,
- $\exists r.A_{s'} \sqsubseteq A_s$ for each $(s, r, s') \in \delta_i$ with $r \in \mathsf{N}_\mathsf{R}^\pm$,
- $A_{s'} \sqcap A \sqsubseteq A_s$ for each $(s, A?, s') \in \delta_i$ with $A \in \mathsf{N}_\mathsf{C}$, and
- $A_{s'} \sqcap A_{s_j} \sqsubseteq A_s$ for each $(s, \langle j \rangle, s') \in \delta_i$, with s_j the initial state of α_j.

We denote by $\mathcal{T}_\mathbf{A}$ the union of all $\mathcal{T}_{\alpha_i}$ with $\alpha_i \in \mathbf{A}$, and define $\mathcal{T}'$ as the union of $\mathcal{T}_\mathbf{A}$ for all atoms $\mathbf{A}_{s_0,F_0}(x, y) \in q$. To obtain the query q' we replace each atom $\mathbf{A}_{s_0,F_0}(x, y)$ by the atom $\alpha'_i(x, y)$, where α_i is the unique automaton in $\mathbf{A}$ with $\{s_0\} \cup F_0 \subseteq S_i$, and α'_i is obtained from α_i by replacing each transition of the form $(s, \langle j \rangle, s') \in \delta_i$ with $(s, A_{s_j}?, s')$, for s_j the initial state of α_j. Note that each α'_i is a standard NFA. We show in the appendix that $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle)$, for every KB $\langle \mathcal{T}, \mathcal{A} \rangle$. □

It follows that in every DL that contains $\mathcal{ELI}$, answering CN2RPQs is no harder than answering plain C2RPQs. From existing upper bounds for C2RPQs (Calvanese, Eiter, and Ortiz 2009; Ortiz, Rudolph, and Simkus 2011), we obtain:

Corollary 4.2. *Answering CN2RPQs is:*

- *in* 2EXP *in combined complexity for all DLs contained in $\mathcal{SHIQ}$, $\mathcal{SHOI}$, $\mathcal{ZIQ}$, or $\mathcal{ZOI}$.*
- *in* EXP *in combined complexity and* P *in data complexity for all DLs contained in Horn-$\mathcal{SHOIQ}$.*

We point out that the 2EXP upper bound for expressive DLs can also be inferred, without using the reduction above, from the existing results for answering C2RPQs in $\mathcal{ZIQ}$ and $\mathcal{ZOI}$ (Calvanese, Eiter, and Ortiz 2009).[4] Indeed, these DLs support regular role expressions as concept constructors, and a nested expression $\langle E \rangle$ in a query can be replaced by a concept $\exists E.\top$ (or by a fresh concept name A_E if the axiom $\exists E.\top \sqsubseteq A_E$ is added to the TBox). Hence, in $\mathcal{ZIQ}$ and $\mathcal{ZOI}$, nested expressions provide no additional expressiveness and CN2RPQs and C2RPQs coincide.

The construction used in Proposition 4.1 also allows us to reduce the evaluation of a N2RPQ to standard reasoning in any DL that contains $\mathcal{ELI}$.

Proposition 4.3. *For every N2RPQ q and every pair of individuals a, b, one can compute in polynomial time an $\mathcal{ELI}$ TBox $\mathcal{T}'$, and a pair of assertions $A_b(b)$ and $A_s(a)$ such that $(a, b) \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ iff $\langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \cup \{A_b(b)\} \rangle \models A_s(a)$, for every DL KB $\langle \mathcal{T}, \mathcal{A} \rangle$.*

From this and existing upper bounds for instance checking in DLs, we easily obtain:

Corollary 4.4. *Answering N2RPQs is in* EXP *in combined complexity for every DL that contains $\mathcal{ELI}$ and is contained in $\mathcal{SHIQ}$, $\mathcal{SHOI}$, $\mathcal{ZIQ}$, or $\mathcal{ZOI}$.*

5 Lower Bounds

The upper bounds we have stated in Section 4 are quite general, and in most cases worst-case optimal.

The 2EXP upper bound stated in the first item of Corollary 4.2 is optimal already for C2RPQs and $\mathcal{ALC}$. Indeed, the 2EXP hardness proof for conjunctive queries in $\mathcal{SH}$ by Eiter et al. (2009) can be adapted to use an $\mathcal{ALC}$ TBox and a C2RPQ. Also the EXP bounds in Corollaries 4.2 and 4.4 are optimal for all DLs that contain $\mathcal{ELI}$, because standard reasoning tasks like satisfiability checking are already EXP-hard in this logic (Baader, Brandt, and Lutz 2008). For the same reasons, the P bound for data complexity in Corollary 4.2 is tight for $\mathcal{EL}$ and its extensions (Calvanese et al. 2006).

However, for the lightweight DLs *DL-Lite*$_\mathcal{R}$ and $\mathcal{EL}$, the best combined complexity lower bounds we have are NL (resp., P) for N2RPQs and PSPACE for CN2RPQs, inherited from the lower bounds for (C)NRPQs (Bienvenu, Ortiz, and Simkus 2013). This leaves a significant gap with respect to the EXP upper bounds in Corollaries 4.2 and 4.4.

We show next that these upper bounds are tight. This is the one of the core technical results of this paper, and probably the most surprising one: already evaluating one N2RPQ in the presence of a *DL-Lite* or $\mathcal{EL}$ TBox is EXP-hard.

Theorem 5.1. *In DL-Lite and $\mathcal{EL}$, N2RPQ answering is* EXP*-hard in combined complexity.*

Proof. We provide a reduction from the word problem for *Alternating Turing Machines (ATMs)* with polynomially

[4] For queries that do not contain inverse roles, that is, (1-way) CRPQs, the same applies to $\mathcal{ZOQ}$ and its sublogics.

bounded space, which is known to be EXP-hard (Chandra, Kozen, and Stockmeyer 1981). An ATM is given as a tuple $M = (\Sigma, S_\exists, S_\forall, \delta, s_{init}, s_{acc}, s_{rej})$, where Σ is an *alphabet*, $S_\exists$ is a set of *existential states*, $S_\forall$ is a set of *universal states*, $\delta \subseteq (S_\exists \cup S_\forall) \times \Sigma \cup \{\mathfrak{b}\} \times (S_\exists \cup S_\forall) \times \Sigma \cup \{\mathfrak{b}\} \times \{-1, 0, +1\}$ is a *transition relation*, $\mathfrak{b}$ is the blank symbol, and $s_{init}, s_{acc}, s_{rej} \in S_\exists$ are the *initial state*, the *acceptance state* and the *rejection state*, respectively.

Consider a word $w \in \Sigma^*$. We can w.l.o.g. assume that $\Sigma = \{0, 1\}$, that M uses only $|w|$ tape cells and that $|w| \geq 1$. Let $m = |w|$, and, for each $1 \leq i \leq m$, let $w(i)$ denote the ith symbol of w. Let $S = S_\exists \cup S_\forall$. We make the following further assumptions:

(i) The initial state is not a final state: $s_{init} \notin \{s_{acc}, s_{rej}\}$.
(ii) Before entering a state s_{acc} or s_{rej}, M writes $\mathfrak{b}$ in all m tape cells.
(iii) There exist functions $\delta_1, \delta_2 : S \times \Sigma \cup \{\mathfrak{b}\} \to S \times \Sigma \cup \{\mathfrak{b}\} \times \{-1, 0, +1\}$ such that $\{\delta_1(s, \sigma), \delta_2(s, \sigma)\} = \{(s', \sigma', d) \mid (s, \sigma, s', \sigma', d) \in \delta\}$ for every $s \in S \setminus \{s_{acc}, s_{rej}\}$ and $\sigma \in \Sigma \cup \{\mathfrak{b}\}$. In other words, non-final states of M give rise to exactly two successor configurations described by the functions δ_1, δ_2.

Note that the machine M can be modified in polynomial time to ensure (i-iii), while preserving the acceptance of w.

We next show how to construct in polynomial time a *DL-Lite* KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and a query q such that M accepts w iff $a \in \mathsf{ans}(q, \mathcal{K})$ (we return to $\mathcal{EL}$ later). The high-level idea underlying the reduction is to use a KB to enforce a tree that contains all possible computations of M on w. The query q selects a computation in this tree and verifies that it corresponds to a proper, error-free, accepting run.

Generating the tree of transitions. First we construct $\mathcal{K}$, which enforces a tree whose edges correspond to the possible transitions of M. More precisely, each edge encodes a transition together with the resulting position of the read/write head of M, and indicates whether the transition is given by δ_1 or δ_2. This is implemented using role names $r_{p,t,i}$, where $p \in \{1, 2\}$, $t \in \delta$, and $0 \leq i \leq m + 1$. To mark the nodes that correspond to the initial (resp., a final) configuration of M, we employ the concept name A_{init} (resp., A_{final}), and we use $A_\exists$ and $A_\forall$ to store the transition type.

We let $\mathcal{A} = \{A_{init}(a), A_\exists(a)\}$, and then we initiate the construction of the tree by including in $\mathcal{T}$ the axiom

$$A_{init} \sqsubseteq \exists r_{p,(s_{init},\sigma,s',\sigma',d),1+d}. \tag{1}$$

for each $\sigma \in \Sigma \cup \{\mathfrak{b}\}$ and $p \in \{1, 2\}$ such that $\delta_p(s_{init}, \sigma) = (s', \sigma', d)$. To generate further transitions, $\mathcal{T}$ contains

$$\exists r^-_{p,(s,\sigma,s',\sigma',d),i} \sqsubseteq \exists r_{p',(s',\sigma^*,s'',\sigma'',d'),i+d'} \tag{2}$$

for each $(s, \sigma, s', \sigma', d) \in \delta$, $1 \leq i \leq m$, $\sigma^* \in \Sigma \cup \{\mathfrak{b}\}$ and $p, p' \in \{1, 2\}$ such that $\delta_{p'}(s', \sigma^*) = (s'', \sigma'', d')$. Note that a transition $t' = (s', \sigma^*, s'', \sigma'', d') \in \delta$ can follow $t = (s, \sigma, s', \sigma', d) \in \delta$ only if σ^* is the symbol written on tape cell i, for i the position of the read/write head after executing t. This is not guaranteed by (2). Instead, we "overestimate" the possible successive transitions, and use the query q to select paths that correspond to a proper computation.

We complete the definition of $\mathcal{T}$ by adding inclusions to label the nodes according to the type of states resulting from transitions. For each $1 \leq i \leq m$, $p \in \{1, 2\}$ and transition $(s, \sigma, s', \sigma', d) \in \delta$, we have the axiom

$$\exists r^-_{p,(s,\sigma,s',\sigma',d),i} \sqsubseteq A_Q, \text{ where}$$

- $A_Q = A_{final}$ if $s' \in \{s_{acc}, s_{rej}\}$,
- $A_Q = A_\exists$, if $s' \in S_\exists \setminus \{s_{acc}, s_{rej}\}$, and
- $A_Q = A_\forall$, if $s' \in S_\forall \setminus \{s_{acc}, s_{rej}\}$.

We turn to the construction of the query q, for which we employ the *n*-NFA representation. We construct an *n*-NFA $\alpha_q = (\mathbf{A}, s, F)$ where $\mathbf{A}$ has $m+1$ automata $\{\alpha_0, \ldots, \alpha_m\}$. Intuitively, the automaton α_0 will be responsible for traversing the tree representing candidate computation paths. At nodes corresponding to the end of a computation path, α_0 launches $\alpha_1, \ldots, \alpha_m$ which "travel" back to the root of the tree and test for the absence of errors along the way. We start by defining the tests $\alpha_1, \ldots, \alpha_m$. Afterwards we define α_0, which selects a set of paths that correspond to a full computation, and launches these tests at the end of each path.

Testing the correctness of a computation path. For each $1 \leq l \leq m$, the automaton $\alpha_l = (S_l, s_l, \delta_l, F_l)$ is built as follows. We let $S_l = \{\sigma_l \mid \sigma \in \Sigma\} \cup \{\mathfrak{b}_l\} \cup \{s'_l\}$. That is, S_l contains a copy of $\Sigma \cup \{\mathfrak{b}\}$ plus the additional state s'_l. We define the initial state as $s_l = \mathfrak{b}_l$ and let $F_l = \{s'_l\}$. Finally, the transition relation δ_l contains the following tuples:

(T1) $(\sigma_l, r^-_{p,(s,\sigma,s',\sigma',d),i}, \sigma_l)$ for all $1 \leq i \leq m$, $p \in \{1, 2\}$, all transitions $(s, \sigma, s', \sigma', d) \in \delta$, and each $\sigma_l \in S_l \setminus \{s'_l\}$ with $l \neq i - d$;
(T2) $(\sigma'_l, r^-_{p,(s,\sigma,s',\sigma',d),i}, \sigma_l)$ for all $1 \leq i \leq m$, $s \in S$ and $p \in \{1, 2\}$ with $\delta_p(s, \sigma) = (s', \sigma', d)$ and $l = i - d$;
(T3) $(\sigma_l, A_{init}?, s'_l)$ for $\sigma = w(l)$.

The working of α_l can be explained as follows. Each state $\sigma_l \in S_l \setminus \{s'_l\}$ corresponds to one of the symbols that may be written in position l of the tape during a run of M. When α_l is launched at some node in a computation tree induced by $\mathcal{K}$, it attempts to travel up to the root node, and the only reason it may fail is when a wrong symbol is written in position l at some point in the computation path. Recall that in each final configuration of M, all symbols are set to the blank symbol, and thus the initial state of α_l is $\mathfrak{b}_l$.

Consider a word $w' \in \mathsf{Roles}^*$ of the form

$$r^-_{p_k,t_k,i_k} \cdots r^-_{p_1,t_1,i_1} \cdot A_{init}? \tag{3}$$

that describes a path from some node in the tree induced by $\mathcal{K}$ up to the root node a. We claim that w' is accepted by every α_l ($1 \leq l \leq m$) just in the case that $t_1, \ldots, t_k$ is a correct sequence of transitions. To see why, first suppose that every α_l accepts w', and let $(pos_0, st_0, tape_0)$ be the tuple with $pos_0 = 1$, $st_0 = s_{init}$ and $tape_0$ contains for each $1 \leq l \leq m$, the symbol σ_l corresponding to the state of α_l when reading A_{init}. Clearly, due to (T3), the tuple $(pos_0, st_0, tape_0)$ describes the initial configuration of M on input w. For $1 \leq j \leq k$, if $t_j = (s, \sigma, s', \sigma', d)$, then we define $(pos_j, st_j, tape_j)$ as follows: $pos_j = i_j$, $st_j = s'$, and $tape_j$ contains for each $1 \leq i \leq m$, the state of α_i when reading $r^-_{p_j,t_j,i_j}$. A simple inductive argument

shows that for every $1 \leq j \leq k$, the tuple $(pos_j, st_j, tape_j)$ describes the configuration of M after applying the transitions $t_1, \ldots, t_j$ from the initial configuration. Indeed, let us assume that $(pos_{j-1}, st_{j-1}, tape_{j-1})$ correctly describes the configuration after executing $t_1, \ldots, t_{j-1}$ and $t_j = (s, \sigma, s', \sigma', d)$. After executing t_j, the read/write head is in position $pos_{j-1} + d$ and the state is s'. Since the only way to enforce an $r^-_{p_j,t_j,i_j}$-edge is via axioms (1) and (2), we must have $pos_j = pos_{j-1} + d$ and $st_j = s'$. It remains to show that $tape_j$ describes the tape contents after executing t_j. Consider some position $1 \leq l \leq m$. There are two cases:

1. $l \neq i_j - d$. In this case, we know that the symbol in position l is not modified by executing t_j. We have to show that $\sigma_l \in tape_{j-1}$ implies $\sigma_l \in tape_j$. This follows from the construction of α_l. In particular, when reading $r_{p_j,t_j,i_j}{}^-$, it must employ a transition from (T1).
2. $l = i_j - d$. In this case, after executing t_j, we must have σ' in position l. We have to show that $\sigma_l \in tape_{j-1}$ implies $\sigma'_l \in tape_j$. This again follows from the construction of α_l. In particular, when reading $r_{p_j,t_j,i_j}{}^-$, there is only one possible transition available in (T2), namely $(\sigma'_l, r_{p_j,t_j,i_j}{}^-, \sigma_l)$.

Conversely, it is easy to see that any word of the form (3) that appears in the tree induced by $\mathcal{K}$ and represents a correct computation path will be accepted by all of the α_l.

Selecting a proper computation. It remains to define α_0, which selects a subtree corresponding to a full candidate computation of M, and then launches the tests defined above at the end of each path. We let $\alpha_0 = (S_0, s_0, \delta_0, F_0)$, where $S_0 = \{s_\downarrow, c_L, c_R, s_\uparrow, s_\updownarrow, s_{test}, s_f\}$, $s_0 = s_\downarrow$, $F_0 = \{s_f\}$, and δ_0 is defined next.

The automaton operates in two main modes: moving down the tree away from the root and moving back up towards the root. Depending on the type of the state of M, in state $s_\downarrow$ the automaton either selects a child node to process next, or chooses to launch the test automata. If the tests are successful, it switches to moving up. To this end, δ_0 has the following transitions:

$$(s_\downarrow, A_\exists?, c_L), (s_\downarrow, A_\exists?, c_R), (s_\downarrow, A_\forall?, c_L),$$
$$(s_\downarrow, A_{final}?, s_{test}), \text{ and } (s_{test}, \langle 1, \ldots, m\rangle, s_\uparrow).$$

The transitions that implement a step down or up are:

- $(c_L, r_{1,t,i}, s_\downarrow)$ for every $1 \leq i \leq m$ and $t \in \delta$,
- $(c_R, r_{2,t,i}, s_\downarrow)$ for every $1 \leq i \leq m$ and $t \in \delta$,
- $(s_\uparrow, r^-_{1,t,i}, s_\updownarrow)$ for every $1 \leq i \leq m$ and $t \in \delta$, and
- $(s_\uparrow, r^-_{2,t,i}, s_\uparrow)$ for every $1 \leq i \leq m$ and $t \in \delta$.

After making a step up from the state $s_\uparrow$ via an $r^-_{1,t,i}$-edge, the automaton enters the state $s_\updownarrow$. Depending on the encountered state of M, the automaton decides either to verify the existence of a computation tree for the alternative transition, to keep moving up, or to accept the word. This is implemented using the following transitions of δ_0:

$$(s_\updownarrow, ?A_\forall, c_R), (s_\updownarrow, ?A_\exists, s_\uparrow), \text{ and } (s_\updownarrow, ?A_{init}, s_f).$$

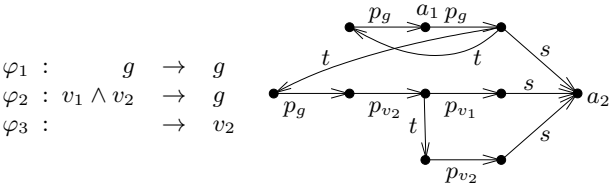


Figure 1: Example ABox in the proof of Theorem 5.2

To conclude the definition of $\alpha_q = (\mathbf{A}, s, F)$, set $s = s_\downarrow$ and $F = \{s_f\}$. Note that α_q has a constant number of states, so it can be converted into an equivalent NRE E_q in polynomial time. The desired query is $q(x, y) = E_q(x, y)$.

The above *DL-Lite* TBox $\mathcal{T}$ can be easily rephrased in $\mathcal{EL}$. Indeed, we simply take a fresh concept name $A_{p,t,i}$ for each role $r_{p,t,i}$, and replace every axiom $C \sqsubseteq \exists r_{p,t,i}$ by $C \sqsubseteq \exists r_{p,t,i}.A_{p,t,i}$ and every axiom $\exists r^-_{p,t,i} \sqsubseteq C$ by $A_{p,t,i} \sqsubseteq C$. $\square$

The above lower bound for answering N2RPQs hinges on the support for existential concepts in the right-hand-side of inclusions. If they are disallowed, then one can find a polynomial time algorithm (Pérez, Arenas, and Gutierrez 2010). However, it was open until now whether the polynomial-time upper bound is optimal. We next prove P-hardness of the problem, already for plain graph databases.

Theorem 5.2. *Given as input an N2RPQ q, a finite interpretation $\mathcal{I}$ and a pair $(o, o') \in \Delta^\mathcal{I} \times \Delta^\mathcal{I}$, it is* P*-hard to check whether $(o, o') \in \mathsf{ans}(q, \mathcal{I})$.*

Proof. To simplify the presentation, we prove the lower bound for a slight reformulation of the problem. In particular, we show P-hardness of deciding $\vec{c} \in \mathsf{ans}(q, \langle \emptyset, \mathcal{A}\rangle)$, where q is an N2RPQ and $\mathcal{A}$ is an ABox with assertions only of the form $A(a)$ or $r(a, b)$, where $A \in \mathsf{N_C}$ and $r \in \mathsf{N_R}$.

We provide a logspace reduction from the classical P-complete problem of checking entailment in propositional definite Horn theories. Assume a set $T = \{\varphi_1, \ldots, \varphi_n\}$ of definite clauses over a set of propositional variables V, where each φ_i is represented as a rule $v_1 \wedge \ldots \wedge v_m \rightarrow v_{m+1}$.

Given a variable $g \in V$, we define an ABox $\mathcal{A}$, an N2RPQ q, and tuple (a_1, a_2) such that $T \models g$ iff $(a_1, a_2) \in \mathsf{ans}(q, \langle \emptyset, \mathcal{A}\rangle)$. We may assume w.l.o.g. that $\varphi_1 = g \rightarrow g$. We define the desired ABox as $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, using the role names s, t, and p_v, where $v \in V$. The ABox $\mathcal{A}_1$ simply encodes T and contains for every $\varphi_i = v_1 \wedge \ldots \wedge v_m \rightarrow v_{m+1}$, the following assertions:

$$p_{v_{m+1}}(e^i_{m+1}, e^i_m), \ldots, p_{v_1}(e^i_1, e^i_0), s(e^i_0, f).$$

The ABox $\mathcal{A}_2$ links variables in rule bodies with their occurrences in rule heads. For every pair of rules $\varphi_i = v_1 \wedge \ldots \wedge v_m \rightarrow v_{m+1}$ and $\varphi_j = w_1 \wedge \ldots \wedge w_n \rightarrow w_{n+1}$, and each $1 \leq l \leq m$ with $v_l = w_{n+1}$, it contains the assertion $t(e^i_{l-1}, e^j_{n+1})$. See Figure 1 for an example.

The existence of a proof tree for g, which can be limited to depth $|V|$, is expressed using the query $q(x, y) = E_{|V|}(x, y)$, with $E_1, E_2, \ldots, E_{|V|}$ defined inductively:

$$E_1 = \bigcup_{v \in V} (p_v \cdot t \cdot p_v) \cdot s$$

$$E_i = \bigcup_{v \in V} \big(p_v \cdot t \cdot p_v\big) \cdot \big(\langle E_{i-1} \rangle \cdot \bigcup_{v \in V} p_v\big)^* \cdot s \qquad (i > 1)$$

Finally, we let $a_1 = e_1^1$ and $a_2 = f$. □

6 Concrete Approach for Horn DLs

Our complexity results so far leave a gap for the data complexity of the *DL-Lite* family: we inherit NL-hardness from plain RPQs, but we only have the P upper bound stemming from Proposition 4.1. In this section, we close this gap by providing an NL upper bound.

This section has an additional goal. We recall that the upper bounds in Corollaries 4.2 and 4.4 rely on reductions to answering (C)2RPQs in extensions of $\mathcal{ELI}$, like Horn-$\mathcal{SHOIQ}$, $\mathcal{ZIQ}$, and $\mathcal{ZOI}$. Unfortunately, known algorithms for C2RPQ answering in these logics use automata-theoretic techniques that are best-case exponential and not considered suitable for implementation. Hence, we want to provide a direct algorithm that may serve as a basis for practicable techniques. To this end, we take an existing algorithm for answering C2RPQs in $\mathcal{ELH}$ and *DL-Lite*$_\mathcal{R}$ due to Bienvenu et al. (2013) and show how it can be extended to handle CN2RPQs and $\mathcal{ELHI}_\bot$ KBs.

For presenting the algorithm in this section, it will be useful to first recall the canonical model property of $\mathcal{ELHI}_\bot$.

Canonical Models

We say that an $\mathcal{ELHI}_\bot$ TBox $\mathcal{T}$ is in *normal form* if all of its concept inclusions are of one of the following forms:

$$A \sqsubseteq \bot \quad A \sqsubseteq \exists r.B \quad \top \sqsubseteq A \quad B_1 \sqcap B_2 \sqsubseteq A \quad \exists r.B \sqsubseteq A$$

with $A, B, B_1, B_2 \in \mathsf{N_C}$ and $r \in \mathsf{N_R^\pm}$.

By introducing fresh concept names to stand for complex concepts, every TBox $\mathcal{T}$ can be transformed in polynomial time into a TBox $\mathcal{T}'$ in normal form that is a model-conservative extension of $\mathcal{T}$. Hence, in what follows, we assume that $\mathcal{ELHI}_\bot$ TBoxes are in normal form.

The domain of the canonical model $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ of a consistent KB $\langle \mathcal{T}, \mathcal{A} \rangle$ consists of all sequences $ar_1C_1 \dots r_nC_n$ ($n \geq 0$) such that:

- $a \in \mathsf{Ind}(\mathcal{A})$ and $r_i \in \mathsf{N_R^\pm}$ for each $1 \leq i \leq n$;
- each C_i is a finite conjunction of concept names;
- if $n \geq 1$, then $\mathcal{T}, \mathcal{A} \models (\exists r_1.C_1)(a)$;
- for $1 \leq i < n$, $\mathcal{T} \models C_i \sqsubseteq \exists r_{i+1}.C_{i+1}$.

For an $o \in \Delta^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} \setminus \mathsf{Ind}(\mathcal{A})$, we use $\mathsf{tail}(o)$ to denote its final concept. The interpretation $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ is then defined as follows:

$$\begin{aligned}
a^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} &= a \text{ for all } a \in \mathsf{Ind}(\mathcal{A}) \\
A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} &= \{a \in \mathsf{Ind}(\mathcal{A}) \mid \mathcal{T}, \mathcal{A} \models A(a)\} \\
&\quad \cup \{o \in \Delta^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} \setminus \mathsf{Ind}(\mathcal{A}) \mid \mathcal{T} \models \mathsf{tail}(o) \sqsubseteq A\} \\
p^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} &= \{(a,b) \mid p(a,b) \in \mathcal{A}\} \cup \\
&\quad \{(o_1,o_2) \mid o_2 = o_1 r\, C \text{ and } \mathcal{T} \models r \sqsubseteq p\} \cup \\
&\quad \{(o_2,o_1) \mid o_2 = o_1 r\, C \text{ and } \mathcal{T} \models r \sqsubseteq p^-\}
\end{aligned}$$

Observe that $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ is composed of a core part containing the individuals from $\mathcal{A}$ and an *anonymous part* consisting of (possibly infinite) trees rooted at the ABox individuals. We use $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_o$ to denote the restriction of $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to domain elements having o as a prefix.

It is well-known that the canonical model of a consistent $\mathcal{ELHI}_\bot$ KB $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ can be homomorphically embedded into any model of $\langle \mathcal{T}, \mathcal{A} \rangle$. Since CN2RPQs are preserved under homomorphisms, we have:

Lemma 6.1. *For every consistent $\mathcal{ELHI}_\bot$ KB $\langle \mathcal{T}, \mathcal{A} \rangle$, CN2RPQ q, and tuple $\vec{a}$ of individuals: $\vec{a} \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ if and only if $\vec{a} \in \mathsf{ans}(q, \mathcal{I}_{\mathcal{T},\mathcal{A}})$.*

Computing Jump and Final Transitions

A crucial component of our algorithm is to compute relevant partial paths in a subtree $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_o$ rooted at an object o in the anonymous part of $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Importantly, we also need to remember which parts of the nested automata that have been partially navigated below o still need to be continued. This will allow us to 'forget' the tree below o.

In what follows, it will be convenient use *runs* to talk about the semantics of *n*-NFAs.

Definition 6.1. *Let $\mathcal{I}$ be an interpretation, and let $(\mathbf{A}, s_0, F_0)$ be an n-NFA. Then a* partial run for $\mathbf{A}$ on $\mathcal{I}$ *is a finite node-labelled tree (T, ℓ) such that every node is labelled with an element from $\Delta^\mathcal{I} \times (\bigcup_i S_i)$ and for each non-leaf node v having label $\ell(v) = (o, s)$ with $s \in S_i$, one of the following holds:*

- *v has a unique child v' with $\ell(v') = (o', s')$, and there exists $(s, \sigma, s') \in \delta_i$ such that $\sigma \in \mathsf{Roles}$ and $(o, o') \in \sigma^\mathcal{I}$;*
- *v has exactly two children v' and v'' with $\ell(v') = (o, s')$ and $\ell(v'') = (o, s'')$, with s'' the initial state of α_j, and there exists a transition $(s, \langle j \rangle, s') \in \delta_i$.*

*If T has root labelled (o_1, s_1) and a leaf node labelled (o_2, s_2) with s_1, s_2 states of the same α_i, then (T, ℓ) is called an (o_1, s_1, o_2, s_2)-*run*, and it is* full *if every leaf label $(o', s') \neq (o_2, s_2)$ is such that $s' \in F_k$ for some k.*

Full runs provide an alternative characterization of the semantics of *n*-NFAs in Definition 3.3.

Fact 6.2. *For every interpretation $\mathcal{I}$, $(o_1, o_2) \in (\mathbf{A}_{s_1,\{s_2\}})^\mathcal{I}$ if and only if there is a full (o_1, s_1, o_2, s_2)-run for $\mathbf{A}$ in $\mathcal{I}$.*

We use partial runs to characterize when an *n*-NFA $\mathbf{A}$ can be partially navigated inside a tree $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_o$ whose root satisfies some conjunction of concepts C. Intuitively, $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ stores pairs s_1, s_2 of states of some $\alpha \in \mathbf{A}$ such that a path from s_1 to s_2 exists, while $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ stores states s_1 for which a path to some final state exists, no matter where the final state is reached. Both $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ and $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ store a set Γ of states s of other automata nested in α, for which a path from s to a final state remains to be found.

Definition 6.2. *Let $\mathcal{T}$ be an $\mathcal{ELHI}_\bot$ TBox in normal form and $(\mathbf{A}, s_0, F_0)$ an n-NFA. The set $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ consists of tuples (C, s_1, s_2, Γ) where C is either $\top$ or a conjunction of concept names from $\mathcal{T}$, s_1 and s_2 are states from $\alpha_i \in \mathbf{A}$, and $\Gamma \subseteq \bigcup_{j>i} S_j$. A tuple (C, s_1, s_2, Γ) belongs to $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ if there exists a partial run (T, ℓ) of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(a)\} \rangle$ that satisfies the following conditions:*

- *the root of T is labelled (a, s_1);*
- *there is a leaf node v with $\ell(v) = (a, s_2)$;*
- *for every leaf node v with $\ell(v) = (o, s) \neq (a, s_2)$, either $s \in F_j$ for some $j > i$, or $o = a$ and $s \in \Gamma$.*

The set $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ contains all tuples (C, s_1, F, Γ) there is a partial run (T, ℓ) of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(a)\}\rangle$ that satisfies the following conditions:

- *the root of T is labelled (a, s_1);*
- *there is a leaf node v with $\ell(v) = (o, s_f)$ and $s_f \in F$;*
- *for every leaf node v with $\ell(v) = (o, s)$, either s is a final state in some α_k, or $o = a$ and $s \in \Gamma$.*

Proposition 6.3. *It can be decided in exponential time if a tuple belongs to $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ or $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$.*

Proof idea. We first show how to use TBox reasoning to decide whether $(C, s_1, s_2, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. For every $\alpha_j \in \mathbf{A}$, we introduce a fresh concept name A_s for each state $s \in S_j$. Intuitively, A_s expresses that there is an outgoing path that starts in s and reaches a final state. If $\{s_1, s_2\} \subseteq S_i$, then we add the following inclusions to $\mathcal{T}$:

- $\top \sqsubseteq A_s$, for every $s \in F_j$ with $j > i$;
- $\exists r.A_{s'} \sqsubseteq A_s$, whenever $(s, r, s') \in \delta_i$ with $r \in \mathsf{N}_\mathsf{R}^\pm$;
- $A_{s'} \sqcap B \sqsubseteq A_s$, whenever $(s, B?, s') \in \delta_i$;
- $A_{s'} \sqcap A_{s''} \sqsubseteq A_s$, whenever $(s, \langle j \rangle, s') \in \delta_i$ and s'' is the initial state of α_j.

Let $\mathcal{T}'$ be the resulting TBox. In the long version, we show that $(C, s_1, s_2, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ iff

$$\mathcal{T}' \models (C \sqcap A_{s_2} \sqcap \prod_{s \in \Gamma} A_s) \sqsubseteq A_{s_1}.$$

To decide if $(C, s_1, F, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$, we must also include in $\mathcal{T}'$ the following inclusions:

- $\top \sqsubseteq A_s$, for every $s \in F$.

We then show that $(C, s_1, F, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ iff

$$\mathcal{T}' \models (C \sqcap \prod_{s \in \Gamma} A_s) \sqsubseteq A_{s_1}.$$

To conclude the proof, we simply note that both problems can be decided in single-exponential time, as TBox reasoning in $\mathcal{ELHI}_\bot$ is known to be EXP-complete. □

Query Rewriting

The core idea of our query answering algorithm is to rewrite a given CN2RPQ q into a set of queries Q such that the answers to q and the union of the answers for all $q' \in Q$ coincide. However, for evaluating each $q' \in Q$, we only need to consider mappings from the variables to the individuals in the core of $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Roughly, a rewriting step makes some assumptions about the query variables that are mapped deepest into the anonymous part and, using the structure of the canonical model, generates a query whose variables are matched one level closer to the core. Note that, even when we assume that no variables are mapped below some element o in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, the satisfaction of the regular paths may require to go below o and back up in different ways. This is handled using jump and final transitions. The query rewriting algorithm is an adaptation of the algorithm for C2RPQs in (Bienvenu, Ortiz, and Simkus 2013), to which the reader may refer for more detailed explanations and examples.

The query rewriting algorithm is presented in Figure 2. In the algorithm, we use atoms of the form $\langle A_{s,F} \rangle(x)$, which are semantically equivalent to $A_{s,F}(x, z)$ for a variable z not occurring anywhere in the query. This alternative notation will spare us additional variables and make the complexity arguments simpler. To slightly simplify the notation, we may write $\mathbf{A}_{s,s'}$ instead of $\mathbf{A}_{s,\{s'\}}$.

The following proposition states the correctness of the rewriting procedure. Its proof follows the ideas outlined above and can be found in the appendix of the long version. Slightly abusing notation, we will also use $\mathsf{Rewrite}(q, \mathcal{T})$ to denote the set all of queries that can be obtained by an execution of the rewriting algorithm on q and $\mathcal{T}$.

Proposition 6.4. *Let $\langle \mathcal{T}, \mathcal{A} \rangle$ be an $\mathcal{ELHI}_\bot$ KB and $q(\vec{x})$ a C2NRPQ. Then $\vec{a} \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ iff there exists $q' \in \mathsf{Rewrite}(q, \mathcal{T})$ and a match π for q' in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ such that $\pi(\vec{x}) = \vec{a}$ and $\pi(y) \in \mathsf{Ind}(\mathcal{A})$ for every variable y in q'.*

We note that the query rewriting does not introduce fresh terms. Moreover, it employs an at most quadratic number of linearly sized *n*-NFAs, obtained from the *n*-NFAs of the input query. Thus, the size of each $q' \in \mathsf{Rewrite}(q, \mathcal{T})$ is polynomial in the size of q and $\mathcal{T}$. Given that all the employed checks in Figure 2 can be done in exponential time (see Proposition 6.3), we obtain the following.

Proposition 6.5. *The set $\mathsf{Rewrite}(q, \mathcal{T})$ can be computed in exponential time in the size of q and $\mathcal{T}$.*

Query Evaluation

In Figure 3, we present an algorithm EvalAtom for evaluating N2RPQs. The idea is similar to the standard nondeterministic algorithm for deciding reachability: we guess a sequence $(c_0, s_0)(c_1, s_1) \cdots (c_m, s_m)$ of individual-state pairs, keeping only two successive elements in memory at any time. Every element (c_{i+1}, s_{i+1}) must be reached from the preceding element (c_i, s_i) by a single normal, jump, or final transition. Moreover, in order to use a jump or final transition, we must ensure that its associated conditions are satisfied. To decide if the current individual belongs to C, we can employ standard reasoning algorithms, but to determine whether an outgoing path exists for one of the states in Γ, we must make a recursive call to EvalAtom. Importantly, these recursive calls involve "lower" automata, and so the depth of recursion is bounded by the number of automata in the N2RPQ (and so is independent of $\mathcal{A}$). It follows that the whole procedure can be implemented in non-deterministic logarithmic space in $|\mathcal{A}|$, if we discount the concept and role membership tests. By exploiting known complexity results for instance checking in *DL-Lite*$_\mathcal{R}$ and $\mathcal{ELHI}_\bot$, we obtain:

Proposition 6.6. EvalAtom *is a sound and complete procedure for N2RPQ evaluation over satisfiable $\mathcal{ELHI}_\bot$ KBs. It can be implemented so as to run in non-deterministic logarithmic space (resp., polynomial time) in the size of the ABox for DL-Lite*$_\mathcal{R}$ *(resp., $\mathcal{ELHI}_\bot$) KBs.*

PROCEDURE Rewrite

Input: CN2RPQ q, $\mathcal{ELHI}_\bot$ TBox $\mathcal{T}$ in normal form

1. Choose either to output q or to continue.
2. Choose a non-empty set $\mathsf{Leaf} \subseteq \mathsf{vars}(q)$ and $y \in \mathsf{Leaf}$. Rename all variables in Leaf to y.
3. Choose a conjunction C of concept names from $\mathcal{T}$ such that $\mathcal{T} \models C \sqsubseteq B$ whenever $B(y)$ is an atom of q. Drop all such atoms from q.
4. For each atom $at \in q$ of the form $\langle \mathbf{A}_{s_0,F} \rangle(t)$ or $\mathbf{A}_{s_0,F}(t,t')$ with $y \in \{t,t'\}$:
 (a) let $\alpha_i \in \mathbf{A}$ be the automaton containing s_0, F
 (b) choose a sequence $s_1, \dots, s_{n-1}$ of distinct states from S_i and some $s_n \in F$
 (c) replace at by the atoms $\mathbf{A}_{s_0,s_1}(t,y)$, $\mathbf{A}_{s_1,s_2}(y,y)$, $\dots$, $\mathbf{A}_{s_{n-2},s_{n-1}}(y,y)$, and
 - $\mathbf{A}_{s_{n-1},s_n}(y,t')$ if $at = \mathbf{A}_{s_0,F}(t,t')$, or
 - $\langle \mathbf{A}_{s_{n-1},s_n} \rangle(y)$ if $at = \langle \mathbf{A}_{s_0,F} \rangle(y)$.
5. For each atom at_j of the form $\mathbf{A}_{s_j,s_{j+1}}(y,y)$ or $\langle \mathbf{A}_{s_j,s_{j+1}} \rangle(y)$ in q, either do nothing, or:
 - Choose some $(C, s_j, s_{j+1}, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ if $at_j = \mathbf{A}_{s_j,s_{j+1}}(y,y)$.
 - Choose some $(C, s_j, \{s_{j+1}\}, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ if $at_j = \langle \mathbf{A}_{s_j,s_{j+1}} \rangle(y)$.
 - Replace at_j by $\{\langle \mathbf{A}_{u,F_k} \rangle(y) \mid u \in \Gamma \cap S_k, \}$.
6. Choose a conjunction D of concept names from $\mathcal{T}$ and $r, r_1, r_2 \in \mathsf{N}_\mathsf{R}^\pm$ such that:
 (a) $\mathcal{T} \models D \sqsubseteq \exists r.C$, $\mathcal{T} \models r \sqsubseteq r_1$, and $\mathcal{T} \models r \sqsubseteq r_2$.
 (b) For each atom $\mathbf{A}_{u,U}(y,t)$ of q with $u \in S_i$, there exists $v \in S_i$ such that $(u, r_1^-, v) \in \delta_i$
 (c) For each atom $\mathbf{A}_{u,U}(t,y)$ of q with $u \in S_i$, there exists $v \in S_i$ and $v' \in U$ with $(v, r_2, v') \in \delta_i$.
 (d) For each atom $\langle \mathbf{A}_{u,U} \rangle(y)$ of q with $u \in S_i$, there exists $v \in S_i$ such that $(u, r_1^-, v) \in \delta_i$.

 For atoms $\mathbf{A}_{u,U}(y,y)$, both (b) and (c) apply.
7. Replace
 - each atom $\mathbf{A}_{u,U}(y,t)$ with $t \neq y$ by $\mathbf{A}_{v,U}(y,t)$,
 - each atom $\mathbf{A}_{u,U}(t,y)$ with $t \neq y$ by $\mathbf{A}_{u,v}(y,t)$,
 - each atom $\mathbf{A}_{u,U}(y,y)$ by atom $\mathbf{A}_{v,v'}(y,y)$, and
 - each atom $\langle \mathbf{A}_{u,U} \rangle(y)$ by atom $\langle \mathbf{A}_{v,U} \rangle(y)$

 with v, v' as in Step 6.
8. Add $A(y)$ to q for each $A \in D$ and return to Step 1.

Figure 2: Query rewriting procedure Rewrite.

We present in Figure 4 the complete procedure EvalQuery for deciding CN2RPQ entailment.

Theorem 6.7. EvalQuery *is a sound and complete procedure for deciding CN2RPQ entailment over $\mathcal{ELHI}_\bot$ KBs. In the case of DL-Lite$_\mathcal{R}$ KBs, it runs in non-deterministic logarithmic space in the size of the ABox.*

Proof idea. Soundness, completeness, and termination of

PROCEDURE EvalAtom

Input: n-NFA $(\mathbf{A}, s_0, F_0)$, $\mathcal{ELHI}_\bot$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ in normal form, $(a,b) \in \mathsf{Ind}(\mathcal{A}) \times (\mathsf{Ind}(\mathcal{A}) \cup \{\mathsf{anon}\})$

1. Let i be such that $s_0 \in S_i$, and set $\mathsf{max} = |\mathcal{A}| \times |S_i| + 1$.
2. Initialize $\mathsf{current} = (a, s_0)$ and $\mathsf{count} = 0$.
3. While $\mathsf{count} < \mathsf{max}$ and $\mathsf{current} \neq (b, s_f)$ for $s_f \in F_0$
 (a) Let $\mathsf{current} = (c, s)$.
 (b) Guess a pair $(d, s') \in (\mathsf{Ind}(\mathcal{A}) \cup \{\mathsf{anon}\}) \times S_i$ such that one of the following holds:
 i. $d \in \mathsf{Ind}(\mathcal{A})$ and there exists $(s, \sigma, s') \in \delta_i$ with $\sigma \in \mathsf{Roles}$ such that $(c,d) \in \sigma^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$
 ii. $d = c$ and $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ contains a tuple (C, s, s', Γ) such that $c \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every $j > i$ and every $u \in \Gamma \cap S_j$,
 $$\mathsf{EvalAtom}((\mathbf{A}, u, F_j), \mathcal{K}, (c, \mathsf{anon})) = \mathtt{yes}$$
 iii. $d = \mathsf{anon}$, $s' \in F_0$, and $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ contains a tuple (C, s, F_0, Γ) such that $c \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every $j > i$ and every $u \in \Gamma \cap S_j$,
 $$\mathsf{EvalAtom}((\mathbf{A}, u, F_j), \mathcal{K}, (c, \mathsf{anon})) = \mathtt{yes}$$
 (c) Set $\mathsf{current} = (d, s')$ and increment count.
4. If $\mathsf{current} = (d, s_f)$ for some $s_f \in F_0$, and either $b = d$ or $b = \mathsf{anon}$, return yes. Else return no.

Figure 3: N2RPQ evaluation procedure EvalAtom.

PROCEDURE EvalQuery

Input: Boolean CN2RPQ q, $\mathcal{ELHI}_\bot$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ in normal form

1. Test whether $\mathcal{K}$ is satisfiable, output yes if not.
2. Set $Q = \mathsf{Rewrite}(q, \mathcal{T})$. Replace all atoms in Q of types $C(a)$, $R(a,b)$ by equivalent atoms of type $\mathbf{A}_{s_0,F_0}(a,b)$.
3. Guess some $q' \in Q$ and an assignment $\vec{a}$ of individuals to the quantified variables $\vec{v}$ in q'
 - Let q'' be obtained by substituting $\vec{a}$ for $\vec{v}$.
 - For every atom $\mathbf{A}_{s_0,F_0}(a,b)$ in q''
 check if $\mathsf{EvalAtom}((\mathbf{A}, s_0, F_0), \mathcal{K}, (a,b)) = \mathtt{yes}$
 - If all checks succeed, return yes.
4. Return no.

Figure 4: CN2RPQ entailment procedure EvalQuery.

EvalQuery follow easily from the corresponding properties of the component procedures Rewrite and EvalAtom (Propositions 6.4, 6.5, and 6.6). In *DL-Lite$_\mathcal{R}$*, KB satisfiability is known to be NL-complete in data complexity. Since the rewriting step is ABox-independent, the size of queries in $\mathcal{Q}$ can be treated as a constant. It follows that the query q' and assignment $\vec{a}$ guessed in Step 3 can be stored in logarithmic space in $|\mathcal{A}|$. By Theorem 6.7, each call to EvalAtom runs in non-deterministic logarithmic space. □

Corollary 6.8. *CN2RPQ entailment over DL-Lite$_\mathcal{R}$ knowl-*

edge bases is NL*-complete in data complexity.*

7 Conclusions and Future Work

We have studied the extension of (C)2RPQs with a nesting construct inspired by XPath, and have characterized the data and combined complexity of answering nested 2RPQs and C2RPQs for a wide range of DLs. The only complexity bound we leave open is whether the coNP lower-bound in data complexity for expressive DLs is tight; indeed, the automata-theoretic approach used to obtain optimal bounds in combined complexity for these logics does not seem to provide the right tool for tight bounds in data complexity.

In light of the surprising jump from P to EXP in the combined complexity of answering N2RPQs in lightweight DLs, a relevant research problem is to identify classes of N2RPQs that exhibit better computational properties. We are also interested in exploring whether the techniques developed in Section 6 can be extended to deal with additional query constructs, such as existential "loop-tests" or forms of role-value maps. Finally, containment of N2RPQs has been studied very recently (Reutter 2013), but only for plain graph databases, so it would be interesting to investigate containment also in the presence of DL constraints.

Acknowledgments. This work has been partially supported by ANR project PAGODA (ANR-12-JS02-007-01), by the EU IP Project FP7-318338 *Scalable End-user Access to Big Data* (Optique), by the FWF projects T515-N23 and P25518-N23, and by the WWTF project ICT12-015.

References

Baader, F.; Calvanese, D.; McGuinness, D.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press.

Baader, F.; Brandt, S.; and Lutz, C. 2008. Pushing the $\mathcal{EL}$ envelope further. In *Proc. of OWLED*.

Barceló Baeza, P. 2013. Querying graph databases. In *Proc. of PODS*, 175–188.

Barceló, P.; Libkin, L.; Lin, A. W.; and Wood, P. T. 2012. Expressive languages for path queries over graph-structured data. *ACM TODS* 37(4):31.

Berglund, A., et al. 2010. XML Path Language (XPath) 2.0 (Second Edition). W3C Recommendation, World Wide Web Consortium. Available at `http://www.w3.org/TR/xpath20`.

Bienvenu, M.; Ortiz, M.; and Simkus, M. 2013. Conjunctive regular path queries in lightweight description logics. In *Proc. of IJCAI*.

Borgida, A., and Brachman, R. J. 2003. Conceptual modeling with description logics. In Baader et al. (2003). chapter 10, 349–372.

Calvanese, D.; De Giacomo, G.; Lenzerini, M.; and Vardi, M. Y. 2000. Containment of conjunctive regular path queries with inverse. In *Proc. of KR*, 176–185.

Calvanese, D.; De Giacomo, G.; Lenzerini, M.; and Vardi, M. Y. 2003. Reasoning on regular path queries. *SIGMOD Record* 32(4):83–92.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2006. Data complexity of query answering in description logics. In *Proc. of KR*, 260–270.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *JAR* 39(3):385–429.

Calvanese, D.; De Giacomo, G.; and Lenzerini, M. 2008. Conjunctive query containment and answering under description logics constraints. *ACM TOCL* 9(3):22.1–22.31.

Calvanese, D.; Eiter, T.; and Ortiz, M. 2009. Regular path queries in expressive description logics with nominals. In *Proc. of IJCAI*, 714–720.

Chandra, A. K.; Kozen, D.; and Stockmeyer, L. J. 1981. Alternation. *JACM* 28(1):114–133.

Consens, M. P., and Mendelzon, A. O. 1990. GraphLog: a visual formalism for real life recursion. In *Proc. of PODS*, 404–416.

Cruz, I. F.; Mendelzon, A. O.; and Wood, P. T. 1987. A graphical query language supporting recursion. In *Proc. of ACM SIGMOD*, 323–330.

Eiter, T.; Lutz, C.; Ortiz, M.; and Simkus, M. 2009. Query answering in description logics with transitive roles. In *Proc. of IJCAI*, 759–764.

Glimm, B.; Lutz, C.; Horrocks, I.; and Sattler, U. 2008. Conjunctive query answering for the description logic $\mathcal{SHIQ}$. *JAIR* 31:157–204.

Harris, S., and Seaborne, A. 2013. SPARQL 1.1 Query Language. W3C Recommendation, World Wide Web Consortium. Available at `http://www.w3.org/TR/sparql11-query`.

Hustadt, U.; Motik, B.; and Sattler, U. 2005. Data complexity of reasoning in very expressive description logics. In *Proc. of IJCAI*.

Krisnadhi, A., and Lutz, C. 2007. Data complexity in the $\mathcal{EL}$ family of description logics. In *Proc. of LPAR*, 333–347.

Lutz, C. 2008. The complexity of conjunctive query answering in expressive description logics. In *Proc. of IJCAR*, volume 5195 of *LNAI*, 179–193. Springer.

Ortiz, M.; Calvanese, D.; and Eiter, T. 2008. Data complexity of query answering in expressive description logics via tableaux. *JAR* 41(1):61–98.

Ortiz, M.; Rudolph, S.; and Simkus, M. 2011. Query answering in the Horn fragments of the description logics $\mathcal{SHOIQ}$ and $\mathcal{SROIQ}$. In *Proc. of IJCAI*, 1039–1044.

Papadimitriou, C. H. 1993. *Computational Complexity*. Addison Wesley.

Pérez, J.; Arenas, M.; and Gutierrez, C. 2010. nSPARQL: A navigational language for RDF. *J. of Web Semantics* 8(4):255–270.

Reutter, J. L. 2013. Containment of nested regular expressions. CoRR Technical Report arXiv:1304.2637, arXiv.org e-Print archive. Available at `http://arxiv.org/abs/1304.2637`.

A Omitted Proofs

Given an n-NFA $(\mathbf{A}, s_0, F_0)$, we define the *level* $\mathsf{lev}(j)$ of $\alpha_j \in \mathbf{A}$ as follows: $\mathsf{lev}(j) = 0$ if the transitions of α_j do not involve any symbol $\langle k \rangle$, and otherwise $\mathsf{lev}(j) = m + 1$ where m is the maximum value of $\mathsf{lev}(k)$ among all $\langle k \rangle$ appearing in a transition of α_j.

Proposition 4.1. *For each CN2RPQ q, one can compute in polynomial time an $\mathcal{ELI}$ TBox $\mathcal{T}'$ and C2RPQ q' such that $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle)$ for every KB $\langle \mathcal{T}, \mathcal{A} \rangle$.*

Proof. Let $\mathcal{T}'$ and q' be defined as in the body of the paper. To complete the argument, we must show $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle)$, for every DL KB $\langle \mathcal{T}, \mathcal{A} \rangle$.

To show the first inclusion $\mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle) \subseteq \mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle)$, it suffices to show that

$$\mathsf{ans}(q, \mathcal{I}) \subseteq \mathsf{ans}(q', \mathcal{I}) \tag{4}$$

for every model $\mathcal{I}$ of $\langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle$. Thus, let $\mathcal{I}$ be an arbitrary model of $\langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle$. To show that Equation 4 holds for $\mathcal{I}$, we prove that for every index j,

$$(\dagger) \text{ if } (o, o') \in \mathbf{A}^{\mathcal{I}}_{s_j, F_j} \text{ for some } o', \text{ then } o \in A^{\mathcal{I}}_{s_j}.$$

Note that it follows from $(\dagger)$ that any sequence $s_0 o_0 s_1 \cdots o_{k-1} s_k o_k$ witnessing $(o, o') \in \mathbf{A}^{\mathcal{I}}_{s_0, F_0}$ also witnesses $(o, o') \in (\alpha'_i)^{\mathcal{I}}$, which yields Equation 4. We prove $(\dagger)$ by induction on the level of α_j. For the base case, assume $\mathsf{lev}(j) = 0$ and $(o, o') \in \mathbf{A}^{\mathcal{I}}_{s_j, F_j}$. Then there are no transitions of the form $(s, \langle j' \rangle, s')$ in δ_j, so there is a sequence $s_0 o_0 s_1 \cdots o_{k-1} s_k o_k$ with $o_0 = o$, $o_k = o'$, $s_0 = s_j$, and $s_k \in F_j$ such that for every $i \in \{1, \ldots, k\}$ there is a transition $(s_{i-1}, \sigma_i, s_i) \in \delta_j$ with $\sigma_i \in \mathsf{Roles}$ and $(o_{i-1}, o_i) \in \sigma_i^{\mathcal{I}}$. Since $\mathcal{I}$ is a model of $\mathcal{T}_{\mathbf{A}}$, we must have $o_k \in A^{\mathcal{I}}_{s_k}$. Using the axioms in $\mathcal{T}_{\mathbf{A}}$ and a simple inductive argument, we can show that $o_i \in A^{\mathcal{I}}_{s_i}$ for all $i \in \{0, \ldots, k\}$, since $(o_{i-1}, o_i) \in \sigma_i^{\mathcal{I}}$. We thus obtain $o \in A^{\mathcal{I}}_{s_j}$. The inductive step is similar, but now the sequence $s_0 o_0 s_1 \cdots o_{k-1} s_k o_k$ may have some $o_{i-1}, s_{i-1}, o_i, s_i$ with $o_{i-1} = o_i$ and $(s_{i-1}, \langle j' \rangle, s_i)$ for some $j' > j$. In this case, by definition there is $o'' \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ such that $(o_i, o'') \in \mathbf{A}^{\mathcal{I}}_{s_{j'}, F_{j'}}$ with $\mathsf{lev}(j') < \mathsf{lev}(j)$, so by the induction hypothesis, $o_{i-1} \in A^{\mathcal{I}}_{s_{j'}}$. This, together with the axioms of $\mathcal{T}_{\mathbf{A}}$, suffices to inductively show $o_i \in A^{\mathcal{I}}_{s_i}$ for all $i \in \{0, \ldots, k\}$, completing the proof of $(\dagger)$.

Next we must show that $\mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle) \subseteq \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$. To this end, for every model $\mathcal{I}$ of $\langle \mathcal{T}, \mathcal{A} \rangle$, we define its extension $\mathcal{I}'$ to the vocabulary of $\mathcal{T}'$ by interpreting the fresh symbols A_s as follows:

$$A_s^{\mathcal{I}'} = \{o \mid \text{ there is } o' \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \text{ such that } (o, o') \in \mathbf{A}^{\mathcal{I}}_{s, F_i}\}$$

where α_i is the unique automaton in $\mathbf{A}$ with $s \in S_i$. Clearly, $\mathcal{I}'$ is a model of $\langle \mathcal{T}, \mathcal{A} \rangle$, and it is not hard to see that it is also a model of $\mathcal{T}'$. Indeed, since the definition of $\mathcal{I}'$ directly reflects the way the axioms simulate the semantics of the automata, a straightforward induction on the level $\mathsf{lev}(j)$ of α_j establishes that $\mathcal{I}' \models \mathcal{T}_{\alpha_j}$ for every index j. Thus, to prove $\mathsf{ans}(q', \langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \rangle) \subseteq \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$, it suffices to show that for every model $\mathcal{I}$ of $\langle \mathcal{T}, \mathcal{A} \rangle$,

$$\mathsf{ans}(q', \mathcal{I}') \subseteq \mathsf{ans}(q, \mathcal{I})$$

with $\mathcal{I}'$ the extension of $\mathcal{I}$ as defined above. Consider some atom $\alpha'_i(x, y)$ of q' and a pair $(o, o') \in (\alpha'_i)^{\mathcal{I}}$. It is easy to see that the sequence $s_0 o_0 s_1 \cdots o_{k-1} s_k o_k$ witnessing $(o, o') \in (\alpha'_i)^{\mathcal{I}'}$ also witnesses $(o, o') \in \mathbf{A}^{\mathcal{I}'}_{s_0, F_0}$ for the n-NFA $\mathbf{A}_{s_0, F_0}$ from which α'_i was obtained, since the automata only differ in the transitions where $\langle j \rangle$ was replaced by $A_{s_j}?$. For every (o_{i-1}, o_i) with $(o_{i-1}, o_i) \in A_{s_j}?^{\mathcal{I}'}$, we have $o_{i-1} = o_i$ and $o_i \in A^{\mathcal{I}'}_{s_j}$. By construction of $\mathcal{I}'$, $o_i \in A^{\mathcal{I}'}_{s_j}$ implies $(o_i, o'') \in \mathbf{A}^{\mathcal{I}}_{s_j, F_j}$ for some o''. Hence $(\alpha'_i)^{\mathcal{I}'} \subseteq \mathbf{A}^{\mathcal{I}}_{s_0, F_0}$ for every atom $\alpha'_i(x, y)$ of q', and so we obtain $\mathsf{ans}(q', \mathcal{I}') \subseteq \mathsf{ans}(q, \mathcal{I})$, as desired. □

Proposition 4.3. *For every N2RPQ q and every pair of individuals a, b, one can compute in polynomial time an $\mathcal{ELI}$ TBox $\mathcal{T}'$, and a pair of assertions $A_b(b)$ and $A_s(a)$ such that $(a, b) \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ iff $\langle \mathcal{T} \cup \mathcal{T}', \mathcal{A} \cup \{A_b(b)\} \rangle \models A_s(a)$, for every DL KB $\langle \mathcal{T}, \mathcal{A} \rangle$.*

Proof. Consider an arbitrary N2RPQ $q(x, y) = E(x, y)$, a pair of individuals a, b, and a knowledge base $\langle \mathcal{T}, \mathcal{A} \rangle$. First we observe that, for every interpretation $\mathcal{I}$, we have $(a, b) \in q^{\mathcal{I}}$ iff $(a, a) \in \langle E \cdot \{b\}? \rangle^{\mathcal{I}}$. Using this observation, the claim follows almost directly from the proof of Proposition 4.1. We take the latter query, and eliminate the symbol $\{b\}$. That is, we add to the ABox the assertion $A_b(b)$, and take the query $q' = \langle E \cdot A_b? \rangle$. Let $\mathcal{A}' = \mathcal{A} \cup \{A_b(b)\}$. We know that $(a, b) \in \mathsf{ans}(q, \langle \mathcal{T}, \mathcal{A} \rangle)$ iff $(a, a) \in \mathsf{ans}(q', \langle \mathcal{T}, \mathcal{A}' \rangle)$.

Let $\mathbf{A}_{s_0, F_0}$ be the n-NFA representation of $\langle E \cdot A_b? \rangle$, and let $\mathcal{T}'$ be $\mathcal{T}_{\mathbf{A}}$ as defined in the proof of Proposition 4.1, except that for the "main" automaton α_i containing states $\{s_0\} \cup F_0$, we use axioms $\top \sqsubseteq A_f$ with $f \in F_0$ rather than $f \in F_i$ (since we may have $F_0 \neq F_i$). We show that $(a, a) \in \mathsf{ans}(q', \langle \mathcal{T}, \mathcal{A}' \rangle)$ iff $\langle \mathcal{T}', \mathcal{A}' \rangle \models A_{s_0}(a)$. For the only if direction, assume $(a, a) \in \mathsf{ans}(q', \langle \mathcal{T}, \mathcal{A}' \rangle)$ and take a model $\mathcal{I}$ of $\langle \mathcal{T}', \mathcal{A}' \rangle$. Since $\mathcal{I} \models \langle \mathcal{T}, \mathcal{A}' \rangle$, we have $(a^{\mathcal{I}}, a^{\mathcal{I}}) \in \mathbf{A}^{\mathcal{I}}_{s_0, F_0}$, and it follows from the proof of Proposition 4.1 that $a^{\mathcal{I}} \in A^{\mathcal{I}}_{s_0}$ as desired. For the other direction, we assume $\langle \mathcal{T}', \mathcal{A}' \rangle \models A_{s_0}(a)$ and consider a model $\mathcal{I}$ of $\langle \mathcal{T}, \mathcal{A}' \rangle$. We have shown in the proof of Proposition 4.1 that $\mathcal{I}$ can be extended to a model $\mathcal{I}'$ of $\langle \mathcal{T}', \mathcal{A}' \rangle$ in such a way that $a^{\mathcal{I}'} \in A^{\mathcal{I}'}_{s_0}$ implies $(a^{\mathcal{I}'}, o) \in \mathbf{A}^{\mathcal{I}'}_{s_0, F_0}$ for some o. Moreover, since $\mathbf{A}_{s_0, F_0}$ is equivalent to an expression of the form $\langle E' \rangle$, $(a^{\mathcal{I}'}, o) \in \mathbf{A}^{\mathcal{I}'}_{s_0, F_0}$ implies $o = a^{\mathcal{I}'}$ and $(a^{\mathcal{I}'}, a^{\mathcal{I}'}) \in \mathbf{A}^{\mathcal{I}'}_{s_0, F_0}$. Hence $(a, a) \in \mathsf{ans}(q', \langle \mathcal{T}, \mathcal{A}' \rangle)$. □

For the next proofs, we will utilize the following facts about canonical models:

Fact A.1. *Let $\mathcal{T}$ be an $\mathcal{ELHI}_\bot$ TBox, and let C, D be $\mathcal{ELHI}_\bot$ concepts. Then $\mathcal{T} \models D \sqsubseteq E$ iff $a \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ where $\mathcal{A} = \{C(a)\}$.*

Fact A.2. *Let $\langle \mathcal{T}, \mathcal{A} \rangle$ be a consistent $\mathcal{ELHI}_\bot$ KB, and suppose that o_1 and o_2 are such that $o_1 \in A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ iff $o_2 \in A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ for every concept name A. Then the submodels $\mathcal{I}_{\mathcal{T}, \mathcal{A}}|_{o_1}$ and $\mathcal{I}_{\mathcal{T}, \mathcal{A}}|_{o_2}$ are isomorphic.*

Proposition 6.3. *It can be decided in* EXP *whether a tuple belongs to* $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ *or* $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$.

Proof. We give the proof only for $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$, but the statement for $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ can be proven in the same fashion. To complete the proof idea presented in the body of the

paper, we must show that $(C, s_1, s_2, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ iff

$$\mathcal{T}' \models (C \sqcap A_{s_2} \sqcap \prod_{s \in \Gamma} A_s) \sqsubseteq A_{s_1} \tag{5}$$

Recall that the states s_1 and s_2 belong to S_i.

For the first direction, suppose that $(C, s_1, s_2, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. By Definition 6.2, there exists a partial run (T, ℓ) of $\mathbf{A}$ in the canonical model $\mathcal{I}$ of $\mathcal{T}$ and $\mathcal{A} = \{C(a)\}$ that satisfies the following conditions:

- the root of T is labelled (a, s_1)
- there is a leaf node labelled (a, s_2)
- for every leaf node labelled $(o, s) \neq (a, s_2)$, either $s \in F_j$ for some $j > i$, or $o = a$ and $s \in \Gamma$

Let $\mathcal{J}$ be the canonical model of $\mathcal{T}'$ and $\mathcal{A}' = \{C(a)\} \cup \{A_{s_2}(a)\} \cup \{A_s(a) \mid s \in \Gamma\}$. Note that $\mathcal{J}$ is also a model of $\langle \mathcal{T}, \mathcal{A} \rangle$, and so by Fact A.2, there must exist a homomorphism $h : \Delta^{\mathcal{I}} \to \Delta^{\mathcal{J}}$ from $\mathcal{I}$ into $\mathcal{J}$ such that $h(a) = a$. By Fact A.1, to show Equation 5, it suffices to show that $a \in {A_{s_1}}^{\mathcal{J}}$. We obtain $a \in {A_{s_1}}^{\mathcal{J}}$ as a consequence of the following claim and the facts that $h(a) = a$ and the root node is labelled (a, s_1).

Claim 1. For every label (o, s) occurring in T: $h(o) \in {A_s}^{\mathcal{J}}$.

Proof of claim. The proof is by induction of the co-depth of nodes in T. For the base case, suppose that v is a leaf node and $\ell(v) = (o, s)$. If $o = a$, then we must have $s = s_2$ or $s \in \Gamma$, and in both cases, we have $A_s(a) \in \mathcal{A}'$. Since $\mathcal{J}$ is a model of $\mathcal{A}'$, we must have $a \in A_s^{\mathcal{J}}$, which yields $h(o) \in {A_s}^{\mathcal{J}}$ since $h(o) = h(a) = a$. Next suppose that $o \neq a$. Then we know from above that $s \in F_j$ for some $j > i$. It follows that $\mathcal{T}'$ contains the inclusion $\top \sqsubseteq A_s$, and so $h(o) \in A_s^{\mathcal{J}}$ trivially holds.

For the induction step, consider a non-leaf node v with $\ell(v) = (o, s)$ and $s \in S_k$, and suppose that the claim has already been shown for the labels of v's children. By Definition 6.1, there are two possibilities for v. The first is that v has a unique child v' with $\ell(v') = (o', s')$ and there is a transition $(s, \sigma, s') \in \delta_k$ such that $\sigma \in \mathsf{Roles}$ and $(o, o') \in \sigma^{\mathcal{I}}$. Since h is a homomorphism, we have $(h(o), h(o')) \in \sigma^{\mathcal{J}}$, and by applying the induction hypothesis, we get $h(o') \in A_{s'}^{\mathcal{J}}$. If $\sigma = r \in \mathsf{N}_{\mathsf{R}}^{\pm}$, then $\mathcal{T}'$ contains the inclusion $\exists r.A_{s'} \sqsubseteq A_s$. Thus, since $\mathcal{J}$ is a model of $\mathcal{T}'$, we must have $h(o) \in A_s^{\mathcal{J}}$. If instead $\sigma = B?$, then we know that $o = o'$ and $o \in B^{\mathcal{I}}$, and using the homomorphism, we get $h(o) = h(o') \in B^{\mathcal{J}}$. The TBox $\mathcal{T}'$ contains the axiom $A_{s'} \sqcap B \sqsubseteq A_s$, and so since $\mathcal{J}$ is a model of $\mathcal{T}'$, we must have $h(o) \in A_s^{\mathcal{J}}$. The second possibility is that the node v has two children v' and v'' with $\ell(v') = (o, s')$ and $\ell(v'') = (o, s'')$ where s'' the initial state of α_j and $(s, \langle j \rangle, s') \in \delta_k$. By the induction hypothesis, we must have $h(o') \in A_{s'}^{\mathcal{J}}$ and $h(o') \in A_{s''}^{\mathcal{J}}$. We also know that $\mathcal{T}'$ contains the inclusion $A_{s'} \sqcap A_{s''} \sqsubseteq A_s$. Since $\mathcal{J}$ is a model of $\mathcal{T}'$, we obtain $h(o) \in A_s^{\mathcal{J}}$, thereby completing the proof of the claim.

For the other direction, suppose that Equation 5 holds. Let $\mathcal{I}$ be the canonical model of $\mathcal{T}$ and $\mathcal{A} = \{C(a)\}$, and let $\mathcal{I}_0$ be the interpretation with the same domain as $\mathcal{I}$ which interprets the concept and role names from $\langle \mathcal{T}, \mathcal{A} \rangle$ exactly as in $\mathcal{I}$, and which additionally interprets each concept name A_s with $s \in \{s_2\} \cup \Gamma$ as $\{a\}$. The interpretation $\mathcal{I}_0$ is a model of $\mathcal{A}' = \{C(a)\} \cup \{A_{s_2}(a)\} \cup \{A_s(a) \mid s \in \Gamma\}$. To extend $\mathcal{I}_0$ to a model of $\mathcal{T}'$, we *chase $\mathcal{I}_0$ with the inclusions in $\mathcal{T}' \setminus \mathcal{T}$*. Formally, we consider an infinite sequence of interpretations $\mathcal{I}_0, \mathcal{I}_1, \mathcal{I}_2, \ldots$ such that $\mathcal{I}_{k+1}$ is obtained from $\mathcal{I}_k$ by selecting an inclusion $D \sqsubseteq A \in \mathcal{T}' \setminus \mathcal{T}$ and an $o \in D^{\mathcal{I}_k} \setminus A^{\mathcal{I}_k}$ and letting $A^{\mathcal{I}_{k+1}} = A^{\mathcal{I}_k} \cup \{o\}$ (and interpreting all other concept and roles names as in $\mathcal{I}_k$). Whenever there are multiple objects o satisfying the condition, we choose an object o having minimal distance from a. Define $\mathcal{J}$ as the interpretation having the same domain as $\mathcal{I}$ and the $\mathcal{I}_k$ and interpreting concept and role names as follows:

$$A^{\mathcal{J}} = \bigcup_{k=0}^{\infty} A^{\mathcal{I}_k} \qquad p^{\mathcal{J}} = \bigcup_{k=0}^{\infty} p^{\mathcal{I}_k}$$

It is not hard to see that $\mathcal{J}$ is a model of $\mathcal{T}'$, and since it is also a model of $\{C(a)\} \cup \{A_{s_2}(a)\} \cup \{A_s(a) \mid s \in \Gamma\}$, we can apply Equation 5 to obtain $a \in A_{s_1}^{\mathcal{J}}$. Moreover, we can find some finite m such that $a \in A_{s_1}^{\mathcal{I}_m}$ and $a \notin A_{s_1}^{\mathcal{I}_{m-1}}$. We now use the sequence $\mathcal{I}_0, \ldots, \mathcal{I}_m$ to construct a partial run of $\mathbf{A}$ on $\mathcal{I}$ that satisfies the conditions of Definition 6.2.

Claim 2. For every $0 \leq k < m$, if $\mathcal{I}_{k+1}$ was obtained from $\mathcal{I}_k$ by applying an inclusion $D \sqsubseteq A_s$ to o, then there is a partial run (T, ℓ) of $\mathbf{A}$ on $\mathcal{I}$ such that:

- the root node is labelled (o, s)
- every leaf node is either labelled by (o', s') with $s' \in F_j$ for some $j > i$, or is labelled (a, s') with $s' \in \{s_2\} \cup \Gamma$.

Proof of claim. The proof is by induction on k. For convenience, throughout we assume that the state s (which corresponds to the concept A_s on the right-hand-side of the inclusion) belongs to S_g. The base case is when $k = 0$. There are four cases to consider depending on the shape of the inclusion that was used to obtain $\mathcal{I}_1$:

- Case 1: $\mathcal{I}_1$ was obtained from $\mathcal{I}_0$ by applying an inclusion $\top \sqsubseteq A_s$ to o, for some $s \in F_j$ with $j > i$. Then the tree with a single node labelled (o, s) defines a partial run with the required conditions.
- Case 2: $\mathcal{I}_1$ was obtained from $\mathcal{I}_0$ by applying $\exists r.A_{s'} \sqsubseteq A_s$ to o. Then we must have $o \in (\exists r.A_{s'})^{\mathcal{I}_0}$, and hence $(o, a) \in r^{\mathcal{I}_0}$, since a is the only object that can belong to $A_{s'}^{\mathcal{I}_0}$. Moreover, since $\exists r.A_{s'} \sqsubseteq A_s$ belongs to $\mathcal{T}$, we must have $(s, r, s') \in \delta_g$. The desired partial run has a root node labelled (o, s) with a single child labelled (a, s').
- Case 3: $\mathcal{I}_1$ was obtained from $\mathcal{I}_0$ by applying $A_{s'} \sqcap B \sqsubseteq A_s$ to o. Then we must have $o \in (A_{s'} \sqcap B)^{\mathcal{I}_0}$, hence $o = a$, $s' \in \{s_2\} \cup \Gamma$, and $a \in B^{\mathcal{I}_0} = B^{\mathcal{I}}$. It follows that $(a, a) \in B?^{\mathcal{I}}$, and since $A_{s'} \sqcap B \sqsubseteq A_s$ belongs to $\mathcal{T}'$, we must also have $(s, B?, s') \in \delta_g$. We can thus use the tree with root node labelled (a, s) and child labelled (a, s').
- Case 4: $\mathcal{I}_1$ was obtained from $\mathcal{I}_0$ by applying $A_{s'} \sqcap A_{s''} \sqsubseteq A_s$ to o. Here again we must have $o = a$, and both s' and s'' must belong to $\{s_2\} \cup \Gamma$. The presence of $A_{s'} \sqcap A_{s''} \sqsubseteq A_s$ in $\mathcal{T}'$ yields $(s, \langle j \rangle, s') \in \delta_g$. Thus, we can use the tree with root labelled (a, s) and children labelled (a, s') and (a, s'').

For the induction step, suppose that the claim holds whenever $0 \leq k < p$. There are again four possibilities to consider for $\mathcal{I}_{p+1}$:

- Case 1: $\mathcal{I}_{p+1}$ was obtained from $\mathcal{I}_p$ by applying an inclusion $\top \sqsubseteq A_s$ to o, for some $s \in F_j$ with $j > i$. As above, we can use the tree with a single node labelled (o, s).
- Case 2: $\mathcal{I}_{p+1}$ was obtained from $\mathcal{I}_p$ by applying $\exists r.A_{s'} \sqsubseteq A_s$ to o. Then we must have $o \in (\exists r.A_{s'})^{\mathcal{I}_p}$, and so there is o' with $o' \in A_{s'}^{\mathcal{I}_p}$. Since $\exists r.A_{s'} \sqsubseteq A_s$ belongs to $\mathcal{T}$, we have $(s, r, s') \in \delta_g$. If $o' = a$ and $s' \in \{s_2\} \cup \Gamma$, then we can use the same argument as in the base case. Else, there must be some earlier stage at which o' was added to $A_{s'}$. We can thus apply the induction hypothesis to find a partial run (T, ℓ) that has root labelled (o', s') and is such that every leaf node is either labelled by (o'', s'') with $s'' \in F_j$ for some $j > i$, or is labelled (a, s'') with $s'' \in \{s_2\} \cup \Gamma$. It then suffices to add a new root node labelled (o, s) as a parent of the root of T to obtain a partial run satisfying the desired conditions.
- Case 3: $\mathcal{I}_{p+1}$ was obtained from $\mathcal{I}_p$ by applying $A_{s'} \sqcap B \sqsubseteq A_s$ to o. Then we must have $o \in A_{s'}^{\mathcal{I}_p}$. If $o = a$ and $s' \in \{s_2\} \cup \Gamma$, then we can proceed as in the base case. Otherwise, we can apply the induction hypothesis to find a partial run (T, ℓ) which has root labelled (o, s') and is such that every leaf node is either labelled by (o', s'') with $s'' \in F_j$ for some $j > i$, or is labelled (a, s'') with $s'' \in \{s_2\} \cup \Gamma$. We also know that $o \in \mathcal{B}^{\mathcal{I}_p} = \mathcal{B}^{\mathcal{I}}$, and that $(s, B?, s') \in \delta_g$ (because of the presence of $A_{s'} \sqcap B \sqsubseteq A_s$ in $\mathcal{T}'$). It follows that we can construct the desired partial run by adding a new root node labelled (o, s) as a parent of the root node of (T, ℓ).
- Case 4: $\mathcal{I}_{p+1}$ was obtained from $\mathcal{I}_p$ by applying $A_{s'} \sqcap A_{s''} \sqsubseteq A_s$ to o. Then we must have $(s, \langle j \rangle, s') \in \delta_g$, and s'' must be the initial state of α_j. We also know that $o \in (A_{s'} \sqcap A_{s''})^{\mathcal{I}_p}$, and so either we can proceed as in the base case, or by applying the induction hypothesis, we can find partial runs (T_1, ℓ_1) and (T_2, ℓ_2) which have roots labelled (o, s') and (o, s'') respectively, and which satisfy the other requirements. Then we can obtain a partial run with the desired properties by creating a new root node labelled (o, s) whose children are the root nodes of (T_1, ℓ_1) and (T_2, ℓ_2).

This completes the proof of Claim 2.

To finish the argument, by taking $k = m - 1$, the above claim yields a partial run (T, ℓ) of $\mathbf{A}$ on $\mathcal{I}$ such that:

- the root node is labelled (a, s_1)
- every leaf node is either labelled by (o', s') with $s' \in F_j$ for some $j > i$, or is labelled (a, s') with $s' \in \{s_2\} \cup \Gamma$.

It is easy to see that if a non-root node is labelled (o', s') and $s' \in S_i$ (recall that $s_1, s_2 \in S_i$), then its parent node must be labelled (o, s) for some $s \in S_i$. Since the states in Γ do not belong to S_i, it must be the case that there is some leaf node labelled (a, s_2). We have thus found a partial run satisfying all of the conditions of Definition 6.2, and so we can conclude that $(C, s_1, s_2, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. □

Proposition 6.4. *$\mathcal{T}, \mathcal{A} \models q$ if and only if there a match π for some query $q' \in \mathsf{Rewrite}(q, \mathcal{T})$ in $\mathcal{I}_{\mathcal{A},\mathcal{T}}$ such that $\pi(t) \in \mathsf{Ind}(\mathcal{A})$ for every $t \in \mathsf{terms}(q')$.*

We split this proof into the two lemmas, a first showing correctness of the procedure Rewrite, and a second showing its completeness. In the proofs of these lemmas, it will prove useful to refer to queries that are produced by a single iteration of Rewrite. We thus introduce the set $\mathsf{one\text{-}step}(q, \mathcal{T})$ which contains precisely those queries q' for which there is an execution of $\mathsf{Rewrite}(q, \mathcal{T})$ such that q' is output the first time that the procedure returns to Step 1.

In the proof below, we write $o \in \langle \mathbf{A}_{s,F} \rangle^{\mathcal{I}}$ to mean that $(o, o') \in \mathbf{A}_{s,F}{}^{\mathcal{I}}$ for some $o' \in \Delta^{\mathcal{I}}$. We also use a notion analogous to an (o_1, s_1, o_2, s_2)-run, but that is more convenient for automata of the form $\langle \mathbf{A}_{s_0,F_0} \rangle$. Let (T, ℓ) be a run for $\mathbf{A}$ in $\mathcal{I}$. We call (T, ℓ) an (o, s, F)-*run* if the root has label (o, s) and some leaf has label (o', s_f) for some $o' \in \Delta^{\mathcal{I}}$ and some $s_f \in S_i \cap F$, with α_i the automaton containing the state s. Naturally, we have that $o \in \langle \mathbf{A}_{s_0,F_0} \rangle^{\mathcal{I}}$ iff there is a full (o, s_0, F_0)-run for $\mathbf{A}$ in $\mathcal{I}$.

Lemma A.3. *If $\mathcal{T}, \mathcal{A} \models q'$ for some $q' \in \mathsf{Rewrite}(q, \mathcal{T})$, then $\mathcal{T}, \mathcal{A} \models q$.*

Proof. It is sufficient to show that if $q' \in \mathsf{one\text{-}step}(q, \mathcal{T})$ and $\mathcal{T}, \mathcal{A} \models q'$, then $\mathcal{T}, \mathcal{A} \models q$. Fix a CN2RPQ q and a $\mathcal{ELHI}$ TBox $\mathcal{T}$. Let $q' \in \mathsf{one\text{-}step}(q, \mathcal{T})$ be such that $\mathcal{T}, \mathcal{A} \models q'$, and let π be a match for q' in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

Consider the execution of $\mathsf{Rewrite}(q, \mathcal{T})$ which leads to the query q' being output the first time that the procedure returns to Step 1. Let Leaf be the non-empty subset of $\mathsf{vars}(q)$ which was selected in Step 2, let C be the conjunction of concept names selected in Step 3, and let D and r, r_1, r_2 be the conjunction of concept names and the roles selected in Step 6. Because of Step 8, we know that q' contains an atom $A(y)$ for each $A \in D$, hence $\pi(y) \in D^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\mathcal{T} \models D \sqsubseteq \exists r.C$, there must exist an r-successor e of $\pi(y)$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ with $\mathsf{tail}(e) = C$. We define a mapping $\pi' : \mathsf{terms}(q) \to \Delta^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ by setting $\pi'(t) = e$ for every $t \in \mathsf{Leaf}$ and setting $\pi'(t) = \pi(t)$ for every $t \in \mathsf{terms}(q') \setminus \{y\}$. This mapping is well-defined since every variable in q either belongs to Leaf or appears in q'.

We aim to show that π' is a match for q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. To this end, consider some concept atom $B(t) \in q$. First suppose that $t \in \mathsf{Leaf}$. Then we know that the concept C selected in Step 3 is such that $\mathcal{T} \models C \sqsubseteq B$. We then use the fact that since $t \in \mathsf{Leaf}$, we have $\pi'(t) = e \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. If $t \notin \mathsf{Leaf}$, then $B(t) \in q'$. As π is a match for q', we have $\pi(t) \in B^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Using $\pi'(t) = \pi(t)$, we get $\pi'(t) \in B^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Now consider some atom $at \in q$ of the form $\mathbf{A}_{s,F}(t, t')$ or $\langle \mathbf{A}_{s,F} \rangle(t)$. For $\mathbf{A}_{s,F}(t, t')$, if both $t \notin \mathsf{Leaf}$ and $t' \notin \mathsf{Leaf}$, then it can be verified that $at \in q'$. As π is a match for q' in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, it must be the case that $(\pi(t), \pi(t')) \in \mathbf{A}_{s,F}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\pi'(t) = \pi(t)$ and $\pi'(t) = \pi(t)$, the same holds for π'. Similarly for $\langle \mathbf{A}_{s,F} \rangle(t)$: if $t \notin \mathsf{Leaf}$, then $at \in q'$, and as π is a match for q' in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, it must be the case that $\pi(t) \in \langle \mathbf{A}_{s,F} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Next we consider the more interesting case in which some term from at is in Leaf. That is, either $at = \mathbf{A}_{s,F}(t, t')$ with $\{t, t'\} \cap \mathsf{Leaf} \neq \emptyset$, or $at = \langle \mathbf{A}_{s,F} \rangle(t)$ with $t \in \mathsf{Leaf}$. Then in Step 4, we have a query containing $\mathbf{A}_{s,F}(\sigma(t), \sigma(t'))$ or $\langle \mathbf{A}_{s,F} \rangle(\sigma(t))$, where either $\sigma(t)$ or $\sigma(t')$ is y, with σ defined as follows: for $t'' \in \{t, t'\}$, $\sigma(t'') = t''$ for $t'' \notin \mathsf{Leaf}$ and

$\sigma(t'') = y$ for $t'' \in \mathsf{Leaf}$. It follows that in Step 4, we will select a sequence $s_1, \dots s_{n-1}$ of distinct states and a final state s_n from the $\alpha_i \in \mathbf{A}$ containing s and F, such that $s_n \in F$. If $at = \mathbf{A}_{s,F}(\sigma(t), \sigma(t'))$, we replace it by the atoms $\mathbf{A}_{s,s_1}(\sigma(t), y)$, $\mathbf{A}_{s_1,s_2}(y, y)$, $\dots$, $\mathbf{A}_{s_{n-2},s_{n-1}}(y, y)$, $\mathbf{A}_{s_{n-1},s_n}(y, \sigma(t'))$. If $at = \langle \mathbf{A}_{s,F} \rangle(y)$, we replace it by the atoms $\mathbf{A}_{s,s_1}(\sigma(t), y)$, $\mathbf{A}_{s_1,s_2}(y, y)$, $\dots$, $\mathbf{A}_{s_{n-2},s_{n-1}}(y, y)$, $\langle \mathbf{A}_{s_{n-1},s_n} \rangle(y)$. Let us denote the former set of atoms by $Q_{\mathbf{A}}$, and the latter by $Q_{\langle \mathbf{A} \rangle}$. We now establish the following claim:

Claim 1.

- If π' is a match for $Q_{\mathbf{A}}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, then π' is a match for $\mathbf{A}_{s,F}(t, t')$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.
- If π' is a match for $Q_{\langle \mathbf{A} \rangle}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, then π' is a match for $\langle \mathbf{A}_{s,F} \rangle(t)$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

Proof of claim. For the first item, suppose that π' is a match for the atoms in $Q_{\mathbf{A}}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Then this means that $(\pi'(\sigma(t)), \pi'(y)) \in (\mathbf{A}_{s,s_1})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, $(\pi'(y), \pi'(y)) \in (\mathbf{A}_{s_i,s_{i+1}})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for every $1 \leq i < n-1$, and $(\pi'(y), \pi'(\sigma(t))) \in (\mathbf{A}_{s_{n-1},s_n})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. It suffices to compose the sequences witnessing each of these membership statements to show that $(\pi'(\sigma(t)), \pi'(\sigma(t'))) \in (\mathbf{A}_{s,s_n})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, hence $(\pi'(\sigma(t)), \pi'(\sigma(t'))) \in (\mathbf{A}_{s,F})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Because of the way we defined π' and σ, we have $\pi'(\sigma(t)) = \pi'(t)$ and $\pi'(\sigma(t')) = \pi'(t')$. We thus obtain $(\pi'(t), \pi'(t')) \in (\mathbf{A}_{s,F})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, i.e. π' is a match for $\mathbf{A}_{s,F}(t, t')$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

For $Q_{\langle \mathbf{A} \rangle}$, the proof is similar. If π' is a match for the atoms in $Q_{\langle \mathbf{A} \rangle}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, then $(\pi'(\sigma(t)), \pi'(y)) \in (\mathbf{A}_{s,s_1})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, $(\pi'(y), \pi'(y)) \in (\mathbf{A}_{s_i,s_{i+1}})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for every $1 \leq i < n-1$, and $(\pi'(y), o) \in (\mathbf{A}_{s_{n-1},s_n})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for some object o. By putting together all the sequences witnessing each atom, we obtain $(\pi'(\sigma(t)), o) \in \langle \mathbf{A}_{s,F} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Then since $\pi'(\sigma(t)) = \pi'(y) = e = \pi'(t)$, we obtain $\pi'(t) \in \langle \mathbf{A}_{s,F} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. (*end proof of Claim 1*)

Because of Claim 1, to complete the proof that π' is a match for q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, it is sufficient to show the following:

Claim 2. The following hold for all atoms in $Q_{\mathbf{A}}$ or $Q_{\langle \mathbf{A} \rangle}$:

- $(\pi'(t_1), \pi'(t_2)) \in \mathbf{A}_{s,s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for atoms of the form $\mathbf{A}_{s,s'}(t_1, t_2)$, and
- $\pi'(t_1) \in \langle \mathbf{A}_{s,s'} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for atoms of the form $\langle \mathbf{A}_{s,s'} \rangle(t_1)$.

Proof of claim. First consider an atom $\mathbf{A}_{s,s'}(t_1, t_2)$ that was not removed in Step 5. There are three cases depending on which of t_1 and t_2 equals y. We treat each case separately:

Case 1: $t_1 = y$ and $t_2 \neq y$. It follows that $t_2 = \sigma(t_2)$ and so $\pi'(t_2) = \pi(t_2)$. In Step 7, we will replace $\mathbf{A}_{s,s'}(t_1, t_2)$ with $\mathbf{A}_{s'',s'}(t_1, t_2)$ where $s'' \in S_i$ is such that $(s, r_1^-, s'') \in \delta_i$. The atom $\mathbf{A}_{s'',s'}(t_1, t_2)$ belongs to q', so we know that it is satisfied by π. More precisely, we know that $(\pi(t_1), \pi(t_2)) \in \mathbf{A}_{s'',s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since e is an r-successor of $\pi(y) = \pi(t_1)$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ and $\mathcal{T} \models r \sqsubseteq r_1$, it follows that $(e, \pi(t_1)) \in \mathbf{A}_{s,s''}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Then putting the above together, and using the fact that $\pi'(t_1) = \pi'(y) = e$, we obtain $(\pi'(t_1), \pi'(t_2)) \in \mathbf{A}_{s,s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Case 2: $t_1 \neq y$ and $t_2 = y$. It follows that $t_1 = \sigma(t_1)$ and so $\pi'(t_1) = \pi(t_1)$. In Step 7, we will replace $\mathbf{A}_{s,s'}(t_1, t_2)$ with an atom $\mathbf{A}_{s,s''}(t_1, t_2)$ where $s'' \in S_i$ is such that $(s'', r_2, s') \in \delta$. The atom $\mathbf{A}_{s,s''}(t_1, t_2)$ appears in q', so it must be satisfied by π. We thus have $(\pi(t_1), \pi(t_2)) \in \mathbf{A}_{s,s''}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. We also know that e is an r-successor of $\pi(t_2) = \pi(y)$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ and that $\mathcal{T} \models r \sqsubseteq r_2$. From this, we can infer that $(\pi(t_2), e) \in \mathbf{A}_{s'',s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ By combining the previous assertions and using the fact that $\pi'(t_2) = e$, we can conclude that $(\pi'(t_1), \pi'(t_2)) \in \mathbf{A}_{s,s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Case 3: $t_1 = t_2 = y$. In Step 7, we will replace $\mathbf{A}_{s,s'}(t_1, t_2)$ with an atom $\mathbf{A}_{v,v'}(t_1, t_2)$ where $(s, r_1^-, v) \in \delta$ and $(v, r_2, s') \in \delta$. By combining the arguments used in Cases 1 and 2, we find that $(\pi'(t_1), \pi(y) \in \mathbf{A}_{s,v}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, $(\pi(y), \pi(y)) \in \mathbf{A}_{v,v'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and $(\pi(y), \pi'(t_2)) \in \mathbf{A}_{v',s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, which together yield $(\pi'(t_1), \pi'(t_2)) \in \mathbf{A}_{s,s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Next we consider the case of an atom of the form $\langle \mathbf{A}_{s,s'} \rangle(y)$ that is not replaced in Step 5. Such an atom is replaced in Step 7 by $\langle \mathbf{A}_{s'',s'} \rangle(y)$, where $s'' \in S_i$ is such that $(s, r_1^-, s'') \in \delta$. The atom $\langle \mathbf{A}_{s'',s'} \rangle(y)$ belongs to q', so we know that it is satisfied by π, and hence $\pi(y) \in \langle \mathbf{A}_{s'',s'} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since e is an r-successor of $\pi(y)$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ and $\mathcal{T} \models r \sqsubseteq r_1$ it follows that $(e, \pi(y) \in \mathbf{A}_{s,s''}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. As $\pi'(y) = e$, we obtain $\pi'(y) \in \mathbf{A}_{s,s'}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

We next treat the case of an atom $\mathbf{A}_{s,s'}(t_1, t_2)$ with $t_1 = t_2 = y$ that is dropped in Step 5 when we chose some $(C, s, s', \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. By definition, we know that there is a partial run (T, ℓ) of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(a)\} \rangle$ such that:

- the root of T is labelled (a, s)
- there is a leaf node labelled (a, s')
- for every leaf node v with $\ell(v) = (o, s'') \neq (a, s')$, either s' is a final state, or $o = a$ and $u \in \Gamma$.

We know from above that $\mathsf{tail}(\pi'(y)) = C$ hence $\pi'(y) \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. By Fact A.2, the canonical model of $\langle \mathcal{T}, \{C(a)\} \rangle$ is isomorphic to $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi'(y)}$. It follows that there is a partial run (T', ℓ') of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi'(y)}$ satisfying the same conditions, except with a replaced by $\pi'(y)$. In Step 5, we added to q' every atom $\langle \mathbf{A}_{u,F_k} \rangle(y)$ such that $u \in \Gamma \cap S_k$. We can thus use exactly the same argument as in the previous case to show that for every $u \in \Gamma \cap S_k$, $\pi'(y) \in \langle \mathbf{A}_{u,F_k} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and hence there is a full $(\pi'(y), u, F_k)$-run (T_u, ℓ_u) of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. By attaching to (T', ℓ') the tree (T_u, ℓ_u) at each non-final leaf v with label $(\pi'(y), u)$, we obtain a full $(\pi'(y), s, \pi'(y), s')$-run for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. This shows that $(\pi'(y), \pi'(y)) \in (\mathbf{A}_{s,s'})^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ which yields the desired result given that $t_1 = t_2 = y$.

The proof is analogous for an atom $\langle \mathbf{A}_{s,s'} \rangle(y)$ in $Q'_{\mathbf{A}}$ that is dropped in Step 5 after choosing some $(C, s, \{s'\}, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$. Again, we know that there is a partial run of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(a)\} \rangle$ as in Definition 6.2, and as $\pi'(y) \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, we can find a corresponding partial run (T', ℓ') in $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi'(y)}$. We know that the root of T' is labelled $(\pi'(y), s)$, there is a leaf labelled (o, s'), and for every leaf node v with $\ell(v) = (o', u) \neq (o, s')$, and u not a final state, we have $o = \pi'(y)$ and $u \in \Gamma$. Since the atoms $\langle \mathbf{A}_{u,F_k} \rangle(y)$ are added

in Step 5, we can use the same arguments as above to find for each such $u \in \Gamma \cap S_k$, a full $(\pi'(y), u, F_k)$-run (T_u, ℓ_u) for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. By attaching to (T', ℓ') the tree (T_u, ℓ_u) at each non-final leaf v with label $(\pi'(y), u)$, we obtain a full $(\pi'(y), s, \{s'\})$-run for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, which shows that $\pi'(y) \in \langle \mathbf{A}_{s,s'} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

This concludes the proof of Claim 2, and hence of Lemma A.3. □

Lemma A.4. *Suppose that $\mathcal{T}, \mathcal{A} \models q$ and π is a match for q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ such that $\pi(y) = p \cdot r_0\, C$ and there is no $z \in \mathsf{vars}(q)$ such that $\pi(y)$ is a proper prefix of $\pi(z)$. Then there is a match π' for some $q' \in \mathsf{one\text{-}step}(q, \mathcal{T})$ such that:*

- *$\pi'(t) = \pi(t)$ for every $t \in \mathsf{terms}(q)$ is such that $\pi(t) \neq \pi(y)$, and*
- *$\pi'(t) = p$ for every $t \in \mathsf{terms}(q)$ with $\pi(t) = \pi(y)$.*

Proof. Let π be a match for a CN2RPQ q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ and t be such that $\pi(y) = p \cdot r_0\, C$ and there is no $z \in \mathsf{vars}(q)$ with $\pi(y)$ a proper prefix of $\pi(z)$. We show how to obtain a query $q' \in \mathsf{one\text{-}step}(q, \mathcal{T})$ and match π' with the required properties. In Step 1 of Rewrite, we choose to continue on to Step 2, where we set $\mathsf{Leaf} = \{t \in \mathsf{vars}(q) \mid \pi(t) = \pi(y)\}$. We define a function σ as follows: $\sigma(t) = y$ if $t \in \mathsf{Leaf}$, and $\sigma(t) = t$ otherwise. At the end of Step 2, we have the query q_2 that contains the atoms $\mathbf{A}_{s,F}(\sigma(t), \sigma(t'))$ for each $\mathbf{A}_{s,F}(t, t') \in q$, $B(\sigma(t))$ for each $B(t) \in q$, and $\langle \mathbf{A}_{s,F} \rangle(\sigma(t))$ for each $\langle \mathbf{A}_{s,F} \rangle(t) \in q$. In Step 3, we choose the concept C. Consider some atom $B(y) \in q_2$. We know that there must be some atom $B(t) \in q$ with $t \in \mathsf{Leaf}$. Since π is a match for q, we must have that $\pi(t) = \pi(y) \in B^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\pi(y) = p \cdot r_0\, C$, it follows from the definition of canonical models that $\mathcal{T} \models C \sqsubseteq B$, as required by Step 3.

Next we show how to select a decomposition of atoms in Step 4. First consider an atom $\mathbf{A}_{s_0,F}(t_1, t_2)$ which is present in the query at the start of Step 4, such that $y \in \{t_1, t_2\}$, and let α_ℓ be the automaton in $\mathbf{A}$ that contains s_0 and F. Then we know from above that there is an atom $\mathbf{A}_{s_0,F}(u, u')$ in the original query q such that $t_1 = \sigma(u)$ and $t_2 = \sigma(u')$. Since π is a match for q with $\pi(t_1) = \pi(u)$ and $\pi(t_2) = \pi(u')$, we know that $(\pi(t_1), \pi(t_2)) \in \mathbf{A}_{s,F}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. This must be witnessed by some sequence $w = u_0 o_0 u_1 \cdots o_{g-1} u_g o_g$ of states of α_ℓ and domain objects as in Definition 3.3, which has $o_0 = \pi(t_1)$, $u_0 = s_0$, $o_g = \pi(t_2)$ and $u_g \in F$. We can assume without loss of generality that g is minimal, i.e. we cannot find a sequence $s'_0 o'_0 \cdots s'_{g'} o'_{g'}$ witnessing $(\pi(t_1), \pi(t_2)) \in \mathbf{A}_{s_0,F}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ with $g' < g$. This ensures that in w there do not exist distinct i, j with $u_i = u_j$, $o_i = o_j$. Now let $j_0 = 0$, $j_n = g$, and $j_1 < \ldots < j_{n-1}$ be all of the indices $1 \leq i < g$ such that $o_i = \pi(y)$. Then define the sequence of states $s_1, \ldots, s_n$ by taking $s_i = u_{j_i}$ for $1 \leq i < n$, and $s_n = u_g$. Note that the states $s_1, \ldots, s_{n-1}$ are all pairwise distinct. We can thus choose this sequence of states in Step 4, and replace the atom $\mathbf{A}_{s_0,F}(t_1, t_2) = \mathbf{A}_{s_0,F}(\sigma(u), \sigma(u'))$ with the atoms: $\mathbf{A}_{s_0,s_1}(\sigma(u), y)$, $\mathbf{A}_{s_1,s_2}(y, y), \ldots,$ $\mathbf{A}_{s_{n-2},s_{n-1}}(y, y)$, $\mathbf{A}_{s_{n-1},s_n}(y, \sigma(u'))$. Note that, by the way we selected the states s_i, we know that $(\pi(t_1), \pi(y)) \in \mathbf{A}_{s,s_1}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, $(\pi(y), \pi(y)) \in \mathbf{A}_{s_i,s_{i+1}}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for $1 \leq i < n-1$, and $(\pi(y), \pi(t_2)) \in \mathbf{A}_{s_{n-1},s_n}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. That is, π is a match for all of these atoms. Moreover, for each of the atoms $\mathbf{A}_{s_i,s_{i+1}}(y, y)$, the sequence witnessing $(\pi(y), \pi(y)) \in \mathbf{A}_{s_i,s_{i+1}}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ is either entirely inside or entirely outside $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$.

For an atom $\langle \mathbf{A}_{s_0,F} \rangle(y)$, the decomposition proceeds analogously. Let α_ℓ be the automaton in $\mathbf{A}$ that contains s_0 and F. Then we know from above that there is an atom $\langle \mathbf{A}_{s_0,F} \rangle(u)$ in the original query q such that $\sigma(u) = y$. Since π is a match for q with $\pi(y) = \pi(u)$, we know that $\pi(y) \in \langle \mathbf{A}_{s_0,F} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and we can take a sequence $w = u_0 o_0 u_1 \cdots o_{g-1} u_g o_g$ of states of α_ℓ and domain objects that witnesses $\pi(y) \in \langle \mathbf{A}_{s_0,F} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Note that we must have $u_0 = s_0$, $o_0 = \pi(y)$, and $u_g \in F$. We can suppose without loss of generality that g is minimal, and so there are no indices $i < j$ such that $u_i = u_j$ and $o_i = o_j$. We let $j_1 < \ldots < j_{n-1}$ be all of the indices $1 \leq i < g$ such that $o_i = \pi(y)$, and for Step 4 use the sequence of states $s_1, \ldots, s_n$ with $s_i = u_{j_i}$ for $1 \leq i < n$, and $s_n = u_g$. Similarly to above, we know that, by construction, $(\pi(y), \pi(y)) \in \mathbf{A}_{s_0,s_1}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, $(\pi(y), \pi(y)) \in \mathbf{A}_{s_i,s_{i+1}}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ for $1 \leq i < n-1$, and $\pi(y) \in \langle \mathbf{A}_{s_{n-1},s_n} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. We also know that, for each of these new atoms $\mathbf{A}_{s_i,s_{i+1}}(y, y)$ or $\langle \mathbf{A}_{s_i,s_{i+1}} \rangle(y)$, the sequence as in Definition 3.3 that witnesses $(\pi(y), \pi(y)) \in \mathbf{A}_{s_i,s_{i+1}}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ or $\pi(y) \in \langle \mathbf{A}_{s_i,s_{i+1}} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ is either entirely inside or entirely outside $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$.

In Step 5, we will remove all atoms $\mathbf{A}_{s_i,s_{i+1}}(y, y)$ or $\langle \mathbf{A}_{s_i,s_{i+1}} \rangle(y)$ for which the witnessing sequence lies entirely inside $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$, possibly introducing some new atoms of the form $\langle \mathbf{A}_{u,F_k} \rangle(y)$. First consider an atom $at = \mathbf{A}_{s_i,s_{i+1}}(y, y)$ such that the witnessing sequence $u_{j_i} o_{j_i} \ldots u_{j_{i+1}} o_{j_{i+1}}$ only involves objects from $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$ (recall that $s_i = u_{j_i}$, $s_{i+1} = u_{j_{i+1}}$, and $o_{j_i} = o_{j_{i+1}} = \pi(y)$). Then there must exist a full $(\pi(y), s_i, \pi(y), s_{i+1})$-run (T, ℓ) for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ whose "main path" has the sequence of labels $(\pi(y), s_i) = (o_{j_i}, u_{j_i}), (o_{j_i+1}, u_{j_i+1}), \ldots, (o_{j_{i+1}}, u_{j_{i+1}}) = (\pi(y), s_{i+1})$. Take this (T, ℓ), and in every branch, look for the first occurrence of a node v_p labeled (o', u') with $o' \notin \Delta^{\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}}$. Let v'_p be the parent of v_p. Note that $\ell(v'_p)$ must be of the form $(\pi(y), u)$. Now, let (T', ℓ) be the result of removing all subtrees rooted at some v_p. Observe that (T', ℓ) is a partial $(\pi(y), s_i, \pi(y), s_{i+1})$-run for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$. Let Γ be the set of states u such that there is a leaf labeled $(\pi(y), u)$ and u is not a final state. Since $\pi(y) = p \cdot r_0 C$, by Fact A.2, $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$ is isomorphic to the canonical model of $\langle \mathcal{T}, C(a) \rangle$, hence there is a partial run for $\mathbf{A}$ in the latter interpretation that has the same features. Therefore we know that $(C, s_i, s_{i+1}, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ and in Step 5, we replace $\mathbf{A}_{s_i,s_{i+1}}(y, y)$ by $\{\langle \mathbf{A}_{u,F_k} \rangle(y) \mid u \in \Gamma \cap S_k\}$. Moreover, for each $u \in \Gamma$, the subtree of T rooted at the node v'_p with $\ell(v'_p) = (\pi(y), u)$ (which became a leaf of T'), together with the labeling ℓ, is a full $(\pi(y), u, F_k)$-run for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, hence $\pi(y) \in \langle \mathbf{A}_{u,F_k} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

The proof is similar when we have an atom $\langle \mathbf{A}_{s_i,s_{i+1}} \rangle(t)$ such that the witnessing sequence $u_{j_i} o_{j_i} \ldots u_{j_{i+1}} o_{j_{i+1}}$ for $\pi(y) \in \langle \mathbf{A}_{s_i,s_{i+1}} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ only involves objects from $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$. Then there must exist a full $(\pi(y), s_i, \{s_{i+1}\})$-run (T, ℓ) for $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ whose sequence of labels along the main path is given by $(\pi(y), s_i) = (o_{j_i}, u_{j_i}), (o_{j_i+1}, u_{j_i+1}), \ldots, (o_{j_{i+1}}, u_{j_{i+1}}) = (o_{j_{i+1}}, s_{i+1})$. From the full run (T, ℓ) in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, we obtain a partial run that allows us to conclude the existence of some $(C, s_i, \{s_{i+1}\}, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$, and we can thus replace at by by atoms $\{\langle \mathbf{A}_{u,F_k} \rangle(y) \mid u \in \Gamma \cap S_k\}$. For all of these atoms, we have $\pi(y) \in \langle \mathbf{A}_{u,F_k} \rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

The final choices to be made occur in Step 6, where we must

choose a conjunction of concept names D, roles $r, r_1, r_2 \in \mathsf{N}_\mathsf{R}^\pm$, and states such that conditions (a)-(d) are satisfied. We set $r = r_0$ (recall that $\pi(y) = p \cdot r_0\, C$) and let D be the conjunction of all concept names A such that $p \in A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. The definition of canonical models, together with our normal form for $\mathcal{ELHI}_\bot$ TBoxes, yields $\mathcal{T} \models D \sqsubseteq \exists r.C$. It remains to show that we can find r_1, r_2 and choices of states such that conditions (a) – (d) are verified. Observe that all atoms of the form $\mathbf{A}_{u,U}(y,t)$, $\mathbf{A}_{u,U}(t,y)$ and $\langle \mathbf{A}_{u,U}\rangle(y)$ in q were obtained by a decomposition in Step 4 (this may include original atoms that were decomposed into only one atom), or were added in Step 5 as a result of some chosen Γ. If we have an atom $\mathbf{A}_{u,U}(t_1,t_2)$ with $y \in \{t_1,t_2\}$, it must have been added in Step 4. We have seen that $(\pi(t_1),\pi(t_2)) \in \mathbf{A}_{u,U}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and since the atom was not removed in Step 5, the witnessing sequence, which we will denote $u_0' o_0' \cdots o_{l-1}' u_l' o_l'$, lies outside the tree $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$: it may start or end at $\pi(y)$, but it contains no other domain element from $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_{\pi(y)}$. We also know that $l \geq 1$ (else again we would have removed the atom in Step 5), and so the sequence must pass the parent p of $\pi(y)$. More specifically, if $\sigma(t_1) = y$ and $u \in S_i$, then $o_0' = \pi(y)$, $u_0' = u$, $o_1' = p$, and and $u_1' = v$ for some $v \in S_i$ such that $(u, r_1^-, v) \in \delta_i$ and $\mathcal{T} \models r_0 \sqsubseteq r_1$. Note that this v is such that $(p,\pi(t_2)) \in \mathbf{A}_{v,U}{}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. If $\sigma(t_2) = y$ and $u \in S_i$, then $o_{l-1}' = p$, $o_l' = \pi(y)$, $u_{l-1}' = v'$, and $u_l = v''$ for some $v', v'' \in S_i$ such that $(v', r_2, v'') \in \delta_i$, $v'' \in U$, and $\mathcal{T} \models r_0 \sqsubseteq r_2$. In this case, we have $(\pi(t_1), p) \in \mathbf{A}_{u,v'}{}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and $(p,\pi(t_2)) \in \mathbf{A}_{v',v''}{}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. If $\sigma(t_1) = \sigma(t_2) = y$, then we pass by p both at the beginning and end of the witnessing sequence, and by choosing states v, v' as above, we obtain $(p,p) \in \mathbf{A}_{v,v'}{}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Finally, if we have an atom of the form $\langle \mathbf{A}_{u,U}\rangle(y)$, then we have seen previously that we must have $\pi(y) \in \langle \mathbf{A}_{u,U}\rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Moreover, since the atom was either added in Step 4, but not dropped in Step 5, or was added in Step 5, it must be the case that the witnessing sequence $u_0' o_0' \cdots o_{l-1}' u_l' o_l'$ for $\pi(y) \in \langle \mathbf{A}_{u,U}\rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ is such that $o_0' = \pi(y)$, $u_0' = u$, $o_1' = p$, and $u_1' = v$ for some $v \in S_i$ such that $(u, r_1^-, v) \in \delta_i$ and and $\mathcal{T} \models r_0 \sqsubseteq r_1$. Note that this v is such that $p \in \langle \mathbf{A}_{v,U}\rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. By selecting $r = r_0, r_1, r_2$ and v, v' in this way, conditions (a) – (d) are all verified.

Now let q' be the query we obtain at the end of Step 8 when all non-deterministic choices are made in the manner described above. Note that $\mathsf{terms}(q') \subseteq \mathsf{terms}(q)$. We aim to find a match π' for q' which satisfies the conditions of the lemma. Let π' be the mapping defined as follows:

- $\pi'(u) = \pi(u)$ for every $u \in \mathsf{terms}(q)$ with $\pi(u) \neq \pi(y)$
- $\pi'(u) = p$ for every $u \in \mathsf{terms}(q)$ with $\pi(u) = \pi(y)$

Clearly, π' satisfies the conditions of the lemma. We only need to show that π' is a match.

To show that π' is a match, first take some concept atom $B(u) \in q'$. There are two possibilities. Either $B(u)$ appears in q and $u \notin \mathsf{Leaf}$, or $B(u)$ was introduced in Step 7. In the former case, we know that π satisfies $B(u)$, and since $\pi'(u) = \pi(u)$ (since $u \notin \mathsf{Leaf}$), the same is true of π'. In the latter case, we must have $u = y$ and $B \in D$. As $\pi'(u) = p$ and D was chosen so that $p \in D^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, π' satisfies $B(u)$.

Now consider some atom $\mathbf{A}_{s,F}(t,t') \in q'$ with $y \notin \{t,t'\}$. Then $\mathbf{A}_{s,F}(t,t') \in q$. As π is a match for q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, it must be the case that $(\pi(t),\pi(t')) \in \mathbf{A}_{s,F}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\pi'(t) = \pi(t)$ and $\pi'(t') = \pi(t')$, the same holds for π', and so the atom $\mathbf{A}_{s,F}(t,t')$ is satisfied by π'. Similarly, for an atom $\langle \mathbf{A}_{s,F}\rangle(t) \in q'$ with $t \neq y$, we know that $\mathbf{A}_{s,F}(t,t') \in q$. As π is a match for q in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, it must be the case that $\pi(t) \in \langle \mathbf{A}_{s,F}\rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\pi'(t) = \pi(t)$, the same holds for π', and so the atom $\langle \mathbf{A}_{s,F}\rangle(t)$ is satisfied by π'.

Next consider some atom $at \in q'$ of the form $\mathbf{A}_{u',U'}(t_1,t_2)$ or $\langle \mathbf{A}_{u',U'}\rangle(t_1)$ with $y \in \{t_1,t_2\}$. A straightforward examination of the procedure Rewrite shows that there is an atom $\mathbf{A}_{s,F}(t',t'')$ or $\langle \mathbf{A}_{s,F}\rangle(t')$ in q which is replaced in Step 4 by the atoms $\mathbf{A}_{s_0,s_1}(\sigma(t'),y)$, $\mathbf{A}_{s_1,s_2}(y,y)$, $\ldots, \mathbf{A}_{s_{n-2},s_{n-1}}(y,y)$, and either $\mathbf{A}_{s_{n-1},s_n}(y,\sigma(t''))$ or $\langle \mathbf{A}_{s_{n-1},s_n}\rangle(y)$, and one of the latter atoms is then replaced by the atom at in Step 7. We distinguish four cases:

Case 1: at replaces $\mathbf{A}_{u,U}(y,t)$ with $t \neq y$. Then at must have the form $\mathbf{A}_{v,U}(y,t)$, where v is the state which was chosen to ensure condition (b) in Step 6. We recall that v is such that $(p,\pi(t)) \in \mathbf{A}_{v,U}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $t \neq y$, we know that $t \notin \mathsf{Leaf}$, and so $\pi(t) = \pi'(t)$. It follows that $(\pi'(y),\pi'(t)) = (p,\pi'(t)) \in \mathbf{A}_{v,U}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and the atom at is satisfied by π'.

Case 2: at replaces $\mathbf{A}_{u,U}(t,y)$ with $t \neq y$. Then at must have the form $\mathbf{A}_{u,v}(t,y)$, where v is the state which was used in condition 6(c). We showed earlier when examining condition 6(c) that $(\pi(t),p) \in \mathbf{A}_{u,v}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Using the fact that $\pi'(t) = \pi(t)$ and $\pi'(y) = p$, we can infer that $(\pi'(t),\pi'(y)) \in \mathbf{A}_{u,v}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, so π' satisfies the atom at.

Case 3: at replaces $\mathbf{A}_{u,U}(y,y)$. Then at must have the form $\mathbf{A}_{v,v'}(y,y)$, where v is the state from 6(b) and v' is the state from 6(c). We have seen above that $(p,p) \in \mathbf{A}_{v,v'}{}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. Since $\pi'(y) = p$, π' satisfies at.

Case 4: at replaces $\langle \mathbf{A}_{u,U}\rangle(y)$. Then at must have the form $\langle \mathbf{A}_{v,U}\rangle(y)$, where v is the state which was chosen to ensure condition (d) in Step 6. We recall that v is such that $p \in \langle \mathbf{A}_{v,U}\rangle^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. As $\pi'(y) = p$, π' satisfies at.

As we have shown that every atom in q' is satisfied by the mapping π', it follows that π' is a match for q' in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, which completes the proof. □

Proposition 6.6. EvalAtom *is a sound and complete procedure for N2RPQ evaluation over satisfiable $\mathcal{ELHI}_\bot$ KBs. It can be implemented so as to run in non-deterministic logarithmic space (resp. polynomial time) in the size of the ABox for DL-Lite$_\mathcal{R}$ (resp. $\mathcal{ELHI}_\bot$) KBs.*

Proof. Consider an n-NFA $(\mathbf{A}, s_0, F_0)$ and a satisfiable $\mathcal{ELHI}_\bot$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A}\rangle$ in normal form. We prove soundness and completeness of EvalAtom by induction on the level of the automaton in the input n-NFA. By Lemma 6.1 and Fact 6.2, it suffices to show the following:

Claim 1. (Soundness and Completeness)

1. There is a full (a, s_0, b, s_f)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ for some $s_f \in F_0$ if and only if there is an execution of EvalAtom which returns yes when given $(\mathbf{A}, s_0, F_0)$, $\langle \mathcal{T}, \mathcal{A}\rangle$, and (a,b) as input, where $b \in \mathsf{Ind}(\mathcal{A})$.

2. There is a full (a, s_0, F_0)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ if and only if there is an execution of EvalAtom which returns yes when given $(\mathbf{A}, s_0, F_0)$, $\langle \mathcal{T}, \mathcal{A}\rangle$, and (a, anon) as input.

Base Case of Claim 1. Suppose that s_0 belongs to an automaton $\alpha_j \in \mathbf{A}$ with $\mathsf{lev}(j) = 0$.

To show the first direction of statement 1, suppose that there is an execution of EvalAtom which returns yes when given $(\mathbf{A}, s_0, F_0)$, $\langle \mathcal{T}, \mathcal{A}\rangle$, and (a, b) as input, As $b \in \mathsf{Ind}(\mathcal{A})$, we know that in Step 3(b), case (iii) does not apply, and because of the way $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ is defined, we know that if $(C, s, s', \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$, then the set Γ must be empty. It follows that when EvalAtom is called on the n-NFA $(\mathbf{A}, s_0, F_0)$, there will be no recursive calls made to EvalAtom. Let $(c_0, s_0)(c_1, s_1) \dots (c_m, s_m)$ be the sequence of elements in current during the execution of EvalAtom. Note that $c_0 = a$ and since EvalAtom returned yes, we also have $c_m = b$ and $s_m \in F_0$. It thus suffices to show the following:

Claim 2. For every $0 \leq i \leq m$, there exists a full (c_i, s_i, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

Proof of claim. The proof is by induction on i. The base case is when $i = m$, in which case we can use the partial run having a single node labelled (c_m, s_m). Now suppose that the claim holds for all $k < i \leq m$, and consider the case when $i = k$. We know that the pair (c_{k+1}, s_{k+1}) must satisfy either condition (i) or (ii) of Step 3(b). If (i) holds, then there exists a transition $(s_k, \sigma, s_{k+1}) \in \delta_j$ with $\sigma \in \mathsf{Roles}$ such that $(c_k, c_{k+1}) \in \sigma^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. From the induction hypothesis, we can find a full $(c_{k+1}, s_{k+1}, c_m, s_m)$-run (T, ℓ) of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Let (T', ℓ') be the labelled tree obtained by creating a new node labelled (c_k, s_k) and making it the parent of the root of (T, ℓ). It is easily verified that (T', ℓ') is defines a full (c_k, s_k, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Next consider the case in which condition (ii) holds for the pair (c_{k+1}, s_{k+1}). Then $c_k = c_{k+1}$ and there exists a tuple $(C, s_k, s_{k+1}, \emptyset)$ in $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ such that $c_k \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. By Definition 6.2, the presence of $(C, s_k, s_{k+1}, \emptyset)$ in $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ means that there is a partial run (T_C, ℓ_C) of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(e)\}\rangle$ that satisfies the following conditions:

- the root of T_C is labelled (e, s_k)
- there is a unique leaf node v with $\ell_C(v) = (e, s_{k+1})$

Since $c_k \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, it follows from Fact A.2 that there is a homomorphism h from the canonical model of $\langle \mathcal{T}, \{C(e)\}\rangle$ to $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ with $h(e) = c_k$. Let (T'_C, ℓ'_C) be obtained by replacing every label (o, s) in T_C by $(h(o), s)$. Using the fact that h is a homomorphism, one can show that (T'_C, ℓ'_C) defines a partial run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ that satisfies the following conditions:

- the root of T'_C is labelled (c_k, s_k)
- there is a unique leaf node v with $\ell'_C(v) = (c_k, s_{k+1})$

Now let (T, ℓ) be the labelled tree obtained from (T'_C, ℓ'_C) by replacing the unique leaf node labelled (c_k, s_{k+1}) by the tree (T_{k+1}, ℓ_{k+1}). It follows from the properties of the component runs that (T, ℓ) is a full (c_k, s_k, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. (*end proof of Claim 2*)

The first direction of statement 2 is proved similarly. We let $(c_0, s_0)(c_1, s_1) \dots (c_m, s_m)$ be the sequence of elements in current during the execution of EvalAtom. If $c_n \in \mathsf{Ind}(\mathcal{A})$, then we can use the same argument as above to construct a full (a, s_0, b, s_m)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Consider next the case in which $c_m = \mathsf{anon}$. Then for (c_m, s_m) to be selected, the conditions in 3(b)(iii) must have satisfied, and so there exists some tuple $(C, s_{m-1}, F_0, \emptyset) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ such that $c_{m-1} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$. It follows that there is a partial run (T, ℓ) for $\mathbf{A}$ on the canonical model $\mathcal{J}$ of $\mathcal{T}$ and $\{C(d)\}$ which consists of a single path and whose sequence ω of labels from root to leaf begins with (c_{m-1}, s_{m-1}) and ends by (o, s_f), for some $o \in \Delta^{\mathcal{J}}$ and $s_f \in F_0$. As $c_{m-1} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, it follows from Fact A.2 that there is a homomorphism h from $\mathcal{J}$ to $\mathcal{I}_{\mathcal{T},\mathcal{A}|_{c_{m-1}}}$ with $h(d) = c_{m-1}$. Let ω' be the sequence obtained by replacing each label (o', s') in ω by $(h(o'), s')$. We then construct a labelled tree consisting of a single path and whose sequence of labels (from root to leaf) is $(c_0, s_0)(c_1, s_1) \dots (c_{m-2}, s_{m-2})\omega'$. It can be verified that this labelled tree defines a full (a, s_0, F_0)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

To prove the second direction of statement 1, suppose that (T, ℓ) is a full (a, s_0, b, s_f)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, where $s_f \in F_0$. Note that because $\mathsf{lev}(j) = 0$, the tree T consists of a single path. Let $(o_0, s_0) \dots (o_m, s_m)$ be the sequence of labels from root to leaf. We have that $o_0 = a$, $o_m = b$, and $s_m = s_f$. Moreover, we know that for each $i \in \{1, \dots, m\}$, there is a transition $(s_{i-1}, \sigma_i, s_i) \in \delta_j$ such that $\sigma_i \in \mathsf{Roles}$ and $(o_{i-1}, o_i) \in \sigma_i^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ (here we use the fact that there are no transitions in α_j of type $\langle k\rangle$). We can suppose without loss of generality that there do not exist $0 \leq i, i' \leq m$ such that $i \neq i'$, $s_i = s_{i'}$, and $o_i = o_{i'}$ (indeed, this condition can always be ensured by deleting a portion of the path). Now let $i_1 < \dots < i_p$ be all of the indices i such that $o_i \in \mathsf{Ind}(\mathcal{A})$ (note that $i_1 = 0$ and $i_p = m$).

Claim 3. For every $1 < l \leq p$, one of the following holds:

1. $(o_{i_{l-1}}, o_{i_l}) \in {\sigma_{i_l}}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, or
2. $o_{i_l} = o_{i_{l-1}}$ and there exists a tuple $(C, s_{i_{l-1}}, s_{i_l}, \emptyset) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ such that $o_{i_{l-1}} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$.

Proof of claim. If $i_l = i_{l-1} + 1$, then we immediately obtain $(o_{i_{l-1}}, o_{i_l}) \in {\sigma_{i_l}}^{\mathcal{I}}$. If $i_l \neq i_{l-1} + 1$, then $o_{i_{l-1}+1} \notin \mathsf{Ind}(\mathcal{A})$. We can infer from this that all the objects $o_{i_{l-1}+1}, \dots, o_{i_l-1}$ must be descendants of $o_{i_l} = o_{i_{l-1}}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Now let C be the conjunction of all concept names A such that $o_{i_{l-1}} \in A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and let $\mathcal{J}$ be the canonical model of $\mathcal{T}$ and $\{C(d)\}$. Then by Fact A.2, there must exist a homomorphism h from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to $\mathcal{J}$ with $h(o_{i_{l-1}}) = d$. Consider the labelled tree (T', ℓ') obtained from (T, ℓ) by (1) making the node labelled $(o_{i_{l-1}}, s_{i_{l-1}})$ the root and the node labelled (o_{i_l}, s_{i_l}) the unique leaf, and then (2) replacing every node label (o', s') by $(h(o'), s')$. Using the fact that h is a homomorphism, one can show that (T', ℓ') is a partial run for $\mathbf{A}$ in $\mathcal{J}$. Moreover, since $o_{i_l} = o_{i_{l-1}}$ and $h(o_{i_{l-1}}) = d$, (T', ℓ') is a $(d, s_{i_{l-1}}, d, s_{i_l})$-run of $\mathbf{A}$ on $\mathcal{J}$. Hence, by Definition 6.2, $(C, s_{i_{l-1}}, s_{i_l}, \emptyset) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. (*end proof of Claim 3*)

Now consider the sequence of pairs $(o_{i_1}, s_{i_1}) \dots (o_{i_p}, s_{i_p})$. Recall that $(o_{i_1}, s_{i_1}) = (a, s_0)$, $(o_{i_p}, s_{i_p}) = (b, s_f)$, $s_f \in F_0$, and by construction $p \leq |\mathcal{A}| \times |S_j|$. Therefore, using Claim 3, we can show that the execution of EvalAtom which guesses this sequence of elements in current during Step 3 will return yes.

The proof of the second direction of statement 2 is broadly similar. We let (T, ℓ) be a full (a, s_0, F_0)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, and we suppose w.l.o.g. that (T, ℓ) is a shortest such run. Let

$(a, s_0) = (o_0, s_0)(o_1, s_1) \dots (o_m, s_m)$ be the sequence of labels in T from root to leaf. Note that we must have $s_m \in F_0$. If $o_m \in \mathsf{Ind}(\mathcal{A})$, then we can use the same proof as above to show that there is a successful execution of EvalAtom. If $o_m \notin \mathsf{Ind}(\mathcal{A})$, then let $i_1 < \dots < i_p$ be all those indices i such that $o_i \in \mathsf{Ind}(\mathcal{A})$. Note that $i_p < m$ since $o_m \notin \mathsf{Ind}(\mathcal{A})$ and that $o_{i_p+1}, \dots, o_m$ must also belong to $\mathcal{I}_{\mathcal{T},\mathcal{A}|_{o_{i_p}}}$. We let C be the conjunction of all concept names A such that $o_{i_p} \in A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and let $\mathcal{J}$ be the canonical model of $\mathcal{T}$ and $\{C(d)\}$. Because of Fact A.2, we know that there exists a homomorphism h from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to $\mathcal{J}$ with $h(o_{i_p}) = d$. Consider the labelled tree (T', ℓ') obtained from (T, ℓ) by (1) making the node labelled (o_{i_p}, s_{i_p}) the root and then (2) replacing every node label (o', s') by $(h(o'), s')$. Using the fact that h is a homomorphism, one can show that (T', ℓ') is a (d, s_{i_p}, F_0) -run of $\mathbf{A}$ on $\mathcal{J}$. We thus have $(C, s_{i_p}, F_0, \emptyset) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ by Definition 6.2. Using this fact together with Claim 3, we can show that the execution of EvalAtom which guesses the sequence $(o_{i_1}, s_{i_1}) \dots (o_{i_p}, s_{i_p})$ of elements in current will return yes. This completes the proof of the base case.

Induction Step of Claim 1. Assume that Claim 1 holds whenever the initial s_0 and final states F_0 come from automaton having level at most g, and let us now consider the case in which s_0 and F_0 belong to α_j whose level is $g + 1$.

To show the first direction of statement 1, suppose that there is an execution of EvalAtom which returns yes when given $(\mathbf{A}, s_0, F_0)$, $\langle \mathcal{T}, \mathcal{A} \rangle$, and (a, b) as input. Let $(c_0, s_0)(c_1, s_1) \dots (c_m, s_m)$ be the sequence of elements placed in current during the top-level execution of EvalAtom. Note that we must have $c_0 = a$ and since EvalAtom returned yes, we must also have $c_m = b$ and $s_m \in F_0$. It thus suffices to show the following:

Claim 4. For every $0 \leq i \leq m$, there exists a full (c_i, s_i, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

Proof of claim. The proof is by induction on i. The base case is when $i = m$, in which case it suffices to consider the tree having a single node labelled (c_m, s_m). Now suppose that the claim holds for all $k < i \leq m$, and consider the case when $i = k$. We know that the pair (c_{k+1}, s_{k+1}) satisfies either condition (i) or (ii) of Step 3(b). If (i) holds, then we can construct a full (c_k, s_k, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, in the same manner as in the base case. Next suppose that it is condition (ii) that holds for the pair (c_{k+1}, s_{k+1}). Then $c_k = c_{k+1}$ and there exists a tuple $(C, s_k, s_{k+1}, \Gamma)$ in $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ such that $c_k \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every state $u \in \Gamma \cap S_l$,

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_l), \mathcal{K}, (c_k, \mathsf{anon})) = \mathtt{yes} \tag{6}$$

By Definition 6.2, the presence of $(C, s_k, s_{k+1}, \Gamma)$ in $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ means that there is a partial run (T_C, ℓ_C) of $\mathbf{A}$ in the canonical model of $\langle \mathcal{T}, \{C(e)\} \rangle$ that satisfies the following conditions:

- the root of T_C is labelled (e, s_k)
- there is a leaf node v with $\ell_C(v) = (e, s_{k+1})$
- for every leaf node v with $\ell_C(v) = (o, s) \neq (e, s_{k+1})$, either $s \in F_{j'}$ for some $j' > j$, or $o = e$ and $s \in \Gamma$

Since $c_k \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, it follows from Fact A.2 that there is a homomorphism h from the canonical model of $\langle \mathcal{T}, \{C(e)\} \rangle$ to $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ with $h(e) = c_k$. Let (T'_C, ℓ'_C) be obtained by replacing every label (o, s) in T_C by $(h(o), s)$. Using the fact that h is a homomorphism, one can show that (T'_C, ℓ'_C) defines a partial run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ that satisfies the following conditions:

- the root of T'_C is labelled (c_k, s_k)
- there is a leaf node v with $\ell'_C(v) = (c_k, s_{k+1})$
- for every leaf node v with $\ell'_C(v) = (o, s) \neq (c_k, s_{k+1})$, either $s \in F_{j'}$ for some $j' > j$, or $o = c_k$ and $s \in \Gamma$

Next, consider the states in Γ. By Equation 6 and the induction hypothesis for Claim 1, we can find for every $u \in \Gamma \cap S_l$ a full (c_k, u, F_l)-run (T_u, ℓ_u) of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Finally, by applying the induction hypothesis of the present claim, and utilizing the fact that $c_k = c_{k+1}$, we can find a full (c_k, s_{k+1}, c_m, s_m)-run (T_{k+1}, ℓ_{k+1}) of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. It then suffices to plug these different runs together. Specifically, we let (T, ℓ) be the labelled tree obtained from (T'_C, ℓ'_C) by:

- replacing the (unique) leaf node labelled (c_k, s_{k+1}) by the tree (T_{k+1}, ℓ_{k+1})
- replacing each leaf node labelled (c_k, u) with $u \in \Gamma \cap S_l$ by the tree (T_u, ℓ_u)

It follows from the properties of the component runs that (T, ℓ) is a full (c_k, s_k, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. (*end proof of Claim 4*)

The proof of the first direction of statement 1 is similar. Indeed, if the final element (c_m, s_m) in current is such that $c_m \in \mathsf{Ind}(\mathcal{A})$, then by Claim 4, there is a full (a, s_0, c_m, s_m)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, which is also a full (a, s_0, F_0)-run. If $c_m = \mathsf{anon}$, then proof is generally similar, except that instead of Claim 4, we use the following claim:

Claim 5. For every $0 \leq i < m$, there exists a full (c_i, s_i, F_0)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$.

Proof of claim. The base case is $i = m - 1$. As $c_m = \mathsf{anon}$, the condition in Step 3(b)(iii) must have been satisfied. Thus, there exists a tuple $(C, s_{m-1}, F_0, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ such that $c_{m-1} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every $u \in \Gamma \cap S_l$,

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_l), \mathcal{K}, (c_{m-1}, \mathsf{anon})) = \mathtt{yes}$$

As $c_{m-1} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and $(C, s_{m-1}, F_0, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$, we can use the same reasoning as in the proof of Claim 4 to obtain a partial run (T_C, ℓ_C) of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ that satisfies the following conditions:

- the root of T_C is labelled (c_{m-1}, s_{m-1})
- there is a leaf node v with $\ell_C(v) = (o, s_f)$ with $s_f \in F_0$
- for every leaf node v with $\ell_C(v) = (o, s)$ with $s \notin F_0$, either $s \in F_{j'}$ for some $j' > j$, or $o = c_{m-1}$ and $s \in \Gamma$

We can also find, for each $u \in \Gamma \cap S_l$, a full (c_k, u, F_l)-run (T_u, ℓ_u) of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. By replacing each leaf node in (T_C, ℓ_C) that is labelled (c_{m-1}, u) with $u \in \Gamma$ by the tree (T_u, ℓ_u), we obtain a full (c_{m-1}, s_{m-1}, F_0)-run of $\mathbf{A}$ on $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. This establishes the base case. The induction step then proceeds exactly as in the proof of Claim 4. (*end proof of Claim 5*)

To show the second direction of statement 1, let us suppose that (T, ℓ) is a full (a, s_0, b, s_f)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, where $s_f \in F_0$. Let $v_0, \dots, v_m$ be the sequence of nodes in T that begins with the root node labelled (a, s_0) and ends with the unique leaf node labelled (b, s_f). We will use (o_l, s_l) for the label of

node n_l. We can assume without loss of generality that (T, ℓ) is minimal in the sense that there do not exist distinct positions $0 \leq l, l' \leq m$ such that n_l and $n_{l'}$ has the same label. Let $i_1 < ... < i_p$ be all of the indices i such that $o_i \in \mathsf{Ind}(\mathcal{A})$. Note that $i_1 = 0$ and $i_p = m$ since $i_0 = a \in \mathsf{Ind}(\mathcal{A})$ and $i_m = b \in \mathsf{Ind}(\mathcal{A})$.

Claim 6. For every $1 < l \leq p$, one of the following holds:

1. there is some $(s_{i_{l-1}}, \sigma, s_{i_l}) \in \delta_j$ with $\sigma \in \mathsf{Roles}$ such that $(o_{i_{l-1}}, o_{i_l}) \in \sigma_{i_l}^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$
2. $o_{i_l} = o_{i_{l-1}}$ and there exists a tuple $(C, s_{i_{l-1}}, s_{i_l}, \Gamma) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ such that $o_{i_l - 1} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every $u \in \Gamma \cap S_k$,

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_k), \mathcal{K}, (o_{i_{l-1}}, \mathsf{anon})) = \mathtt{yes}$$

Proof of claim. It follows from Definition 6.1 that there are three possibilities:

(a) the node $n_{i_{l-1}}$ has n_{i_l} as its unique child
(b) the node $n_{i_{l-1}}$ has n_{i_l} as one of its two children
(c) the node n_{i_l} is a descendant, but not a child, of $n_{i_{l-1}}$

First consider case (a). Then by Definition 6.1, there must exist a transition $(s_{i_{l-1}}, \sigma, s_{i_l}) \in \delta_j$ such that $\sigma \in \mathsf{Roles}$ and $(o_{i_{l-1}}, o_{i_l}) \in \sigma^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, so the first statement of the claim is satisfied. If (b) holds, then it follows from Definition 6.1 that $(s_{i_{l-1}}, \langle k \rangle, s_{i_l}) \in \delta_j$, $o_{i_l} = o_{i_{l-1}}$, and the other child of $n_{i_{l-1}}$ has label $(o_{i_{l-1}}, u)$, with u the initial state of α_k. Since $(s_{i_{l-1}}, \langle k \rangle, s_{i_l}) \in \delta_j$, we have $(\top, s_{i_{l-1}}, s_{i_l}, \{u\}) \in \mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$. The labelled subtree of (T, ℓ) rooted at $(o_{i_{l-1}}, u)$ is a full $(o_{i_{l-1}}, u, F_k)$-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, and since $k > j$, we have $\mathsf{lev}(k) \leq g$. The induction hypothesis for Claim 1 is thus applicable and yields

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_k), \mathcal{K}, (o_{i_{l-1}}, \mathsf{anon})) = \mathtt{yes}$$

We have thus shown that the second statement of the claim holds. Finally, let us consider case (c). Since n_{i_l} is not a child of $n_{i_{l-1}}$, we must have $o_{i_{l-1}+1} \notin \mathsf{Ind}(\mathcal{A})$. We can infer from this that all the objects $o_{i_{l-1}+1}, \ldots, o_{i_l - 1}$ must be descendants of $o_{i_{l-1}}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, and that $o_{i_{l-1}} = o_{i_l}$. Consider the partial run (T', ℓ') obtained from (T, ℓ) by

1. making $n_{i_{l-1}}$ the new root node
2. for every descendant v of $n_{i_{l-1}}$ such that $\ell(v) = (o_{i_{l-1}}, s)$ and such that there is no v' (different from v and $n_{i_{l-1}}$) with $\ell(v') = (o_{i_{l-1}}, s')$ and occurring along the path from $n_{i_{l-1}}$ to v, make v a leaf node by dropping all children of v

Note that since (T, ℓ) is a full run, there are three types of leaf nodes in (T', ℓ'):

- the node n_{i_l} (on the main path) with label $(o_{i_{l-1}}, s_{i_l})$
- nodes with labels of the form $(o_{i_{l-1}}, s)$
- nodes with labels of the form (o, s_f) with $s_f \in F_k$ for some $k > j$

Moreover, by construction, every label (o, s) appearing in (T', ℓ') is such that $o \in \mathcal{I}_{\mathcal{T},\mathcal{A}}|_{o_{i_{l-1}}}$. Now let C be the conjunction of all concept names A such that $o_{i_{l-1}} \in A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and let $\mathcal{J}$ be the canonical model of $\mathcal{T}$ and $\{C(d)\}$. Then it follows from Fact A.2 there exists a homomorphism h from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to $\mathcal{J}$ with $h(o_{i_{l-1}}) = d$. Consider the labelled tree (T'', ℓ'') obtained from (T', ℓ') by replacing every node label (o, s) by $(h(o), s)$. Using the fact that h is a homomorphism with $h(o_{i_{l-1}}) = d$ and the above description of the leaf nodes in (T', ℓ'), one can show that (T'', ℓ'') is a partial run of $\mathbf{A}$ on $\mathcal{J}$ that satisfies:

- the root of T'' is labelled $(d, s_{i_{l-1}})$
- there is a leaf node labelled (d, s_{i_l})
- for every leaf node v with $\ell(v) = (o, s) \neq (d, s_{i_l})$, either $s \in F_k$ for some $k > j$, or $o = d$

Thus, by Definition 6.2, the tuple $(C, s_{i_{l-1}}, s_{i_l}, \Gamma)$ belongs to $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$, if we let Γ contain all those states u such that there is a leaf node (d, u) with $u \neq s_{i_l}$. To complete the proof, consider some $u \in \Gamma \cap S_k$. Then there is a leaf node of (T'', ℓ'') with label (d, u), hence a leaf node of (T', ℓ') with label $(o_{i_{l-1}}, u)$, which also appears (perhaps not as a leaf) in the original run (T, ℓ). Moreover, since (T, ℓ) is a full run, by taking the subtree of (T, ℓ) rooted at this node, we obtain a full $(o_{i_{l-1}}, u)$-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. As u belongs to α_k, with $k > j$, the induction hypothesis for Claim 1 is applicable, and yields

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_k), \mathcal{K}, (o_{i_{l-1}}, \mathsf{anon})) = \mathtt{yes}$$

We have thus shown that the second statement of the claim holds. (*end proof of Claim 6*)

It follows from Claim 6 and the fact that $p \leq |\mathcal{A}| \times |S_j|$ that there is an execution of EvalAtom which guesses the sequence $(o_{i_1}, s_{i_1}) \ldots (o_{i_p}, s_{i_p})$ of elements in current that returns yes.

The second direction of statement 2 can be proven similarly. We suppose that (T, ℓ) is a full (a, s_0, F_0)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$, and let $v_0, \ldots, v_m$ be the sequence of nodes in T that begins with the root node labelled (a, s_0) and ends with the unique leaf node labelled (o, s) with $s \in S_j$. We use (o_l, s_l) for the label of node n_l. Note that since (T, ℓ) is a full (a, s_0, F_0)-run, we have $s_m \in F_0$. We define the sequence $i_1 < \ldots < i_p$ and make the same minimality assumption on (T, ℓ) to ensure that $p \leq |\mathcal{A}| \times |S_j|$. If $o_m \in \mathsf{Ind}(\mathcal{A})$, we can follow the above proof for statement 1 exactly. If $o_m \notin \mathsf{Ind}(\mathcal{A})$, then we must additionally show that:

Claim 7. There is a tuple $(C, s_{i_p}, F_0, \Gamma) \in \mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ such that $o_{i_p} \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ and for every $u \in \Gamma \cap S_k$,

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_k), \mathcal{K}, (o_{i_{l-1}}, \mathsf{anon})) = \mathtt{yes}$$

Proof of claim. Let (T', ℓ') be the partial run obtained from the run (T, ℓ) by

1. making n_{i_p} the new root node
2. for every descendant v of n_{i_p} such that $\ell(v) = (o_{i_p}, s)$ and such that there is no v' (different from v and n_{i_p}) with $\ell(v') = (o_{i_p}, s')$ and occurring along the path from n_{i_p} to v, make v a leaf node by dropping all of its children

Note that since (T, ℓ) is a full run, there are three types of leaf nodes in (T', ℓ'):

- the node n_m (on the main path) with label (o_m, s_m)
- nodes with labels of the form (o_{i_p}, s)
- nodes with labels of the form (o, s_f) with $s_f \in F_k$ for some $k > j$

Moreover, by construction, every label (o, s) appearing in (T', ℓ') is such that $o \in \mathcal{I}_{\mathcal{T},\mathcal{A}}|_{o_{i_p}}$. Now let C be the conjunction of all concept names A such that $o_{i_p} \in A^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$, and let $\mathcal{J}$ be the canonical model of $\mathcal{T}$ and $\{C(d)\}$. By Fact A.2, there exists a homomorphism h from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to $\mathcal{J}$ with $h(o_{i_p}) = d$. Consider the labelled tree (T'', ℓ'') obtained from (T', ℓ') by replacing every node label (o, s) by $(h(o), s)$. Using the fact that h is a homomorphism with $h(o_{i_p}) = d$ and the above description of the leaf nodes in (T', ℓ'), one can show that (T'', ℓ'') is a partial run of $\mathbf{A}$ on $\mathcal{J}$ that satisfies:

- the root of T'' is labelled (d, s_{i_p})
- there is a leaf node labelled (d, s_m)
- for every leaf node v with $\ell(v) = (o, s) \neq (d, s_m)$, either $s \in F_k$ for some $k > j$, or $o = d$

Then since $s_m \in F_0$, by Definition 6.2, the tuple $(C, s_{i_p}, F_0, \Gamma)$ belongs to $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$, for the Γ containing exactly those states u such that there is a leaf node (d, u) with $u \neq s_m$. To complete the proof, consider some $u \in \Gamma \cap S_k$. Then there is a leaf node of (T'', ℓ'') with label (d, u), and hence a leaf node of (T', ℓ') with label (o_{i_p}, u). By the construction of (T', ℓ'), there is a node labelled (o_{i_p}, u) in the original run (T, ℓ), and since (T, ℓ) is a full run, there is a full (o_{i_p}, u, F_k)-run of $\mathbf{A}$ in $\mathcal{I}_{\mathcal{T},\mathcal{A}}$. Applying the induction hypothesis for Claim 1, we obtain

$$\mathsf{EvalAtom}((\mathbf{A}, u, F_k), \mathcal{K}, (o_{i_p}, \mathsf{anon})) = \mathtt{yes}$$

(*end proof of Claim 7*)

Using Claims 6 and 7 and the fact that $p \leq |\mathcal{A}| \times |S_j|$, one can show that there is an execution of EvalAtom which guesses the sequence $(o_{i_1}, s_{i_1}) \dots (o_{i_p}, s_{i_p})(\mathsf{anon}, s_m)$ of elements in current that returns yes. (*end proof of Claim 1*)

We now show the second part of the proposition, namely that EvalAtom runs in non-deterministic logarithmic space in the size of $\mathcal{A}$. First observe that when the input *n*-NFA is $(\mathbf{A}, s_0, F_0)$, the depth of the recursion cannot exceed $\mathsf{lev}(s_0)$. The latter number is bounded above by the number of automata in $\mathbf{A}$, and hence is independent of the size of $\mathcal{A}$. Next we note that for each call of EvalAtom, there are only three pieces of information that must be stored and whose size depends on $|\mathcal{A}|$: the individual c in current, the "next" individual d, and the value of the counter count. Each of these pieces of information can be stored using logarithmic space in $|\mathcal{A}|$. We then remark that testing whether a tuple belongs to $\mathsf{JumpTrans}(\mathbf{A}, \mathcal{T})$ or $\mathsf{FinalTrans}(\mathbf{A}, \mathcal{T})$ can be done in constant time in $|\mathcal{A}|$, since it is ABox-independent. The only checks that involve $\mathcal{A}$ are testing whether $(c, d) \in \sigma^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ in Step 3(b)(i) and testing whether $c \in C^{\mathcal{I}_{\mathcal{T},\mathcal{A}}}$ in Steps 3(b)(ii) and 3(b)(iii). If $\mathcal{T}$ is formulated in $\mathcal{ELHI}_\bot$, then both checks can be done in polynomial time in $|\mathcal{A}|$, and for *DL-Lite*$_\mathcal{R}$, this can be improved to non-deterministic logarithmic space in $|\mathcal{A}|$. This follows from the known P (resp. NL) data complexity of instance checking in these logics (Hustadt, Motik, and Sattler 2005; Calvanese et al. 2007). Putting this all together, we have that EvalAtom runs in polynomial time in $|\mathcal{A}|$ for $\mathcal{ELHI}_\bot$ KBs, and in non-deterministic logarithmic space in the size of $\mathcal{A}$ for *DL-Lite*$_\mathcal{R}$ KBs. □